The Standard Manual for Engineering Creature Comforts

Volume One — The External Instrument

Michael L. Curzi · 36N9 GENETICS LLC

Contents

Part I — The Circuit **5**

Chapter 1 . Recursion as Circuit, Not Roadblock **6**
The Law . . . 6
The Engineering . . . 10
The Specification . . . 18
The Failure Modes . . . 19
Worked Example . . . 21

Chapter 2 . The Leaky Integrator **25**
The Law . . . 25
The Engineering . . . 27
The Specification . . . 37
The Failure Modes . . . 38
Worked Example . . . 40

Chapter 3 . Loss — Engineering Delta **45**
The Law . . . 45
The Engineering . . . 47
The Specification . . . 56
The Failure Modes . . . 57
Worked Example . . . 60

Chapter 4 . Conversion — Engineering Eta **66**
The Law . . . 66
The Engineering . . . 67
The Specification . . . 77
The Failure Modes . . . 78
Worked Example . . . 80

Chapter 5 . Cost — The Dissipation Budget **84**
The Law . . . 84
The Engineering . . . 85
The Specification . . . 96
The Failure Modes . . . 97

Worked Example 99

Chapter 6 . Headroom — The Master Gauge **102**
The Law 102
The Engineering 103
The Specification 115
The Failure Modes 116
Worked Example 118

Part II — The Geometry of Difference **123**

Chapter 7 . The Effective Count **124**
The Law 124
The Engineering 127
The Specification 136
The Failure Modes 137
Worked Example 139

Chapter 8 . Raising N — The Logarithmic Ceiling **145**
The Law 145
The Engineering 146
The Specification 157
The Failure Modes 158
Worked Example 161

Chapter 9 . Engineering u — Designing Non-Parallelism **167**
The Law 167
The Engineering 170
The Specification 180
The Failure Modes 181
Worked Example 184

Chapter 10 . Crossing and Divergence **190**
The Law 190
The Engineering 193
The Specification 202
The Failure Modes 204
Worked Example 207

Chapter 11 . Pseudo-Diversity and the Parallel Wall **212**
The Law 212
The Engineering 215
The Specification 226
The Failure Modes 228
Worked Example 231

Part III — Engineering Space **235**

Chapter 12 . The Room as Circuit Element **236**

The Law . . . 236
The Engineering . . . 237
The Specification . . . 249
The Failure Modes . . . 253
Worked Example . . . 256

Chapter 13 . Fractal and Parabolic Space **263**
The Law . . . 263
The Engineering . . . 264
The Specification . . . 278
The Failure Modes . . . 279
Worked Example . . . 281

Chapter 14 . Acoustics for Turn-Taking **288**
The Law . . . 288
The Engineering . . . 290
The Specification . . . 310
The Failure Modes . . . 313
Worked Example . . . 315

Chapter 15 . Thresholds, Edges, and Transitions **322**
The Law . . . 322
The Engineering . . . 323
The Specification . . . 334
The Failure Modes . . . 335
Worked Example . . . 336

Chapter 16 . Light, Material, and Thermal **340**
The Law . . . 340
The Engineering . . . 341
The Specification . . . 353
The Failure Modes . . . 355
Worked Example . . . 356

Chapter 17 . The Table With No Head **360**
The Law . . . 360
The Engineering . . . 362
The Specification . . . 372
The Failure Modes . . . 373
Worked Example . . . 374

Part IV — The Sociology **379**

Chapter 18 . Collective Intelligence as an Engineerable Quantity **380**
The Law . . . 380
The Engineering . . . 381
The Specification . . . 391
The Failure Modes . . . 393
Worked Example . . . 394

Chapter 19 . Contagion, Thresholds, and Critical Mass **400**
The Law . . . 400
The Engineering . . . 401
The Specification . . . 410
The Failure Modes . . . 412
Worked Example . . . 413

Chapter 20 . The Silo as Local Minimum **418**
The Law . . . 418
The Engineering . . . 419
The Specification . . . 428
The Failure Modes . . . 429
Worked Example . . . 431

Chapter 21 . Order Without a Centre **438**
The Law . . . 438
The Engineering . . . 439
The Specification . . . 450
The Failure Modes . . . 452
Worked Example — The Ashfield Tool Library . . . 455

Chapter 22 . Individuality as Load-Bearing Structure **461**
The Law . . . 461
The Engineering . . . 463
The Specification . . . 473
The Failure Modes . . . 474
Worked Example — A Nine-Person Review Panel . . . 476

Chapter 23 . Institutions as Storage **481**
The Law . . . 481
The Engineering . . . 484
The Specification . . . 492
The Failure Modes . . . 493
Worked Example — A Technical Standards Body, 1962 to 2026 . . . 496

Part V — Tradecraft **502**

Chapter 24 . The Survey **503**
The Law . . . 503
The Engineering . . . 505
The Specification . . . 519
The Failure Modes . . . 523
Worked Example . . . 525

Chapter 25 . Preregistration and Falsification **532**
The Law . . . 532
The Engineering . . . 534
The Specification . . . 544
The Failure Modes . . . 547

Worked Example 548

Chapter 26 . Intervention Design **554**
The Law 554
The Engineering 556
The Specification 571
The Failure Modes 575
Worked Example 577

Chapter 27 . The Failure Mode Catalog **587**
The Law 587
The Engineering — the catalog 588
The Specification — how to use the catalog 602
The Failure Modes — of diagnosis itself 603
Worked Example — a hospital ward shift handover 604

Chapter 28 . Worked Systems **608**
The Law 608
The Engineering — four systems 609
The Specification 627
The Failure Modes 628
Worked Example — one desk, and the scale-free claim 629

Part I — The Circuit

Chapter 1 . Recursion as Circuit, Not Roadblock

Book One taught an operator to read a comfort gradient — to notice which room, which table, which conversation made the work easier, and to trust that reading as data. This manual takes the next step. It teaches an engineer to *build* the circumstance that generates the gradient and then holds onto it.

The whole manual rests on one equation, and that equation is a recursion. So the first thing this book has to do is repair the reputation of recursion.

Recursion has a bad name in engineering practice, and the bad name was earned in one narrow place: the call stack. A function that calls itself without a base case consumes a frame per call until the frame allocator runs out. The failure is loud, total, and famous enough to have become a metaphor: infinite regress, the mirror facing a mirror, the turtles all the way down, the argument that never bottoms out. When an organisation says "we're going in circles," it means it has found itself inside that metaphor.

That metaphor describes exactly one kind of recursion: **recursion without a loss term.** It is not the kind this manual is about, and it is not the kind that most real systems are.

The storage law of the Interaction Surplus Framework is a recursion *with* a loss term. It reads

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t
```

and the small coefficient `delta` in the first term is the entire difference between a regress and a circuit. **Loss is not the defect in the loop. Loss is the thing that turns the loop into a device that arrives somewhere.**

This chapter proves that claim, shows what goes wrong on each side of it, and lays out the method of the rest of the manual.

The Law

1.1 The two laws this book runs on

Two results from the Interaction Surplus Framework (Curzi, 2026) are used everywhere in this volume. Both are **PROVEN** within their papers. Neither is an empirical claim.

The Surplus Law (Paper A, Theorem 2.1). For two unit state vectors x and y drawn from an ensemble of N categories,

```
F(x,y) = f(u(x,y)),    u(x,y) = 1 − (x . y)^2 = ||x ^ y||^2   in [0,1]

f(u) = ln(1 + (N−1) u)
```

and `f` is **uniquely** characterised by four axioms:

Axiom	Statement	Status
S1	Geometric dependence: `F` depends on the pair only through u	structural

Axiom	Statement	Status
S2	Zero at zero: `f(0) = 0`	structural
S3	Affine effective count: `g(u) = e^{f(u)}` is affine in u	**DESIGN AXIOM — chosen, not derived**
S4	Normalisation: `f(1) = ln N`	structural

S3 is the one that does the real work, and it is a choice. **We flag it every time we lean on it: the affine effective count is a design axiom, not a discovery.** It says we have decided to measure difference by a headcount that grows linearly in the squared sine of the angle between states. Given that decision, the logarithm is forced. The decision itself is ours.

The object S3 defines is the **effective count**:

`N_eff = g(u) = 1 + (N−1) u`

and surplus is `ln N_eff`. **Surplus is the logarithm of a headcount after the headcount has been audited for redundancy.** Ten instruments that agree perfectly have `u = 0`, `N_eff = 1`, and surplus `ln 1 = 0`. They are one instrument that has been purchased ten times.

Derived from the axioms, not assumed: strict monotonicity (Paper A, 3.1), positivity (3.2), strict concavity (3.3), subadditivity on the parameter u (3.4), sharp Lipschitz constant N−1 (4.2), and extremality — `F = 0` if and only if the states are parallel, `F = ln N` if and only if they are orthogonal (4.3).

Range (Paper D, Corollary 2.2): `0 <= F(u) <= ln N`. **The quantity `ln N` is the hard ceiling on surplus per interaction**, and it will appear in every design inequality in this manual.

The Storage Law (Paper E, §3). This is the engine of the book.

`Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,    S_t = F(U_t) = ln(1 + (N−1) U_t)`

Symbol	Meaning	Range
`Q_t`	stored order: what persists — institution, habit, furnished room, trust, runbook	`>= 0`
`S_t`	surplus generated by interaction this period	`0 <= S_t <= ln N`
`delta`	retention loss — the fraction of stored order lost per period without maintenance	`[0,1]`
`eta`	conversion efficiency — interaction into *durable* structure	`>= 0`
`C_t`	cost, dissipation, overhead	`>= 0`

Admissibility: `C_t <= (1−delta) Q_t + eta S_t` guarantees `Q_{t+1} >= 0`. A period that violates it is not a period in which the model predicts a negative stock; it is a period in which something has broken that the linear model was never built to represent.

Exact solution (Theorem 3.1):

`Q_t = (1−delta)^t Q_0 + sum_{k=0}^{t−1} (1−delta)^{t−1−k} (eta S_k − C_k)`

One-step growth (Corollary 3.2): `Q_{t+1} > Q_t` if and only if `eta S_t > delta Q_t + C_t`.

Uniform bound (Theorem 3.3): `Q_t <= max{Q_0, eta ln N / delta}`. **If `delta = 0`, the bound degrades to `Q_t <= Q_0 + t eta ln N`** — linear in `t`, unbounded, no ceiling.

Sustainability ceiling (Corollary 3.4): a nondecreasing step requires `C_t <= eta ln N − delta Q_t`. If `C_t > eta ln N − delta Q_t`, **no choice of interaction geometry can prevent order from decreasing.** This is the single most important sociological statement in the book, and it is proved, not asserted.

Steady state (Theorem 3.5): with constant `U == u*`, constant `C == c*`, and `delta > 0`,

`Q_inf = (eta F(u*) − c*) / delta`

and if `eta F(u*) < c*` the equilibrium is negative — meaning the regime is unsustainable or the linear model must be replaced. The theorem does not hide this failure mode, and neither will we.

Headroom (Paper E, §6.4), the master gauge of this manual:

`h_t := eta ln N − delta Q_t − C_t`

with `h_t < 0` unsustainable in this model, `h_t = 0` maintenance only, `h_t > 0` growth available with any requested increment `r_t <= h_t`. The maximum one-step increment (Corollary 4.2) is exactly `r_max = eta ln N − delta Q_t − C_t`.

1.2 The central identity — recursion as circuit

The storage law is a leaky integrator. In discrete time it is exactly a first-order low-pass filter, a charge accumulator with a leak:

```
Q          <–>  charge on a capacitor
delta      <–>  leakage conductance
eta S      <–>  current in
C          <–>  load draw
(1 − delta)<–>  feedback gain / pole location
Q_inf      <–>  steady−state charge
```

This is not an analogy chosen for flavour. It is the same difference equation, and every standard result about first-order discrete filters — pole, time constant, DC gain, step response, settling time — applies without modification. Chapter 2 develops the correspondence in full.

Recursion here is not a stack that overflows. It is an accumulator that integrates.

1.3 Banach — why iteration arrives

Write the update as a map on the real line. Hold `S_t = S` and `C_t = C` constant and define

`T(Q) = (1 − delta) Q + eta S − C`

For any two states `Q` and `Q'`,

`|T(Q) − T(Q')| = (1 − delta) |Q − Q'|`

so T is Lipschitz with modulus `k = 1 − delta`. If `delta > 0` then `k < 1` and T is a **contraction**. The Banach fixed point theorem then gives three things at once, on a complete metric space, with no further hypotheses:

1. **Existence.** A fixed point `Q*` exists.
2. **Uniqueness.** There is exactly one.
3. **Global convergence.** Iterating T from *any* starting point converges to `Q*`, with the geometric error bound `|Q_t − Q*| <= k^t |Q_0 − Q*|`.

Solving `Q* = (1−delta)Q* + eta S − C` gives `Q* = (eta S − C)/delta`, which is Paper E's Theorem 3.5. **Book One's Node 10 FIXEDPOINT is Paper E Theorem 3.5.** They are the same statement in two registers: the operator's experience of *arriving somewhere and staying there* is the mathematician's unique attracting fixed point of a contraction.

Book One's Node 09 RECURSE is the transient term `(1−delta)^t Q_0` — the memory of where you started, decaying geometrically while the driven term builds. Node 09 and Node 10 are not two phenomena. They are the two terms of Theorem 3.1.

And the whole apparatus turns on the single inequality `k = 1 − delta < 1`, which is true if and only if `delta > 0`.

Set `delta = 0` and Banach does not apply. There is no unique fixed point, there is no convergence, and there is no arrival. The map becomes an isometry: every point is a fixed point, distances are preserved forever, and the system remembers its initial condition with weight exactly 1 for all time.

That is the thesis of this chapter, stated as tersely as it can be stated.

1.4 The honesty clauses

These are repeated wherever the law is used, and they are repeated here because this is the chapter that introduces the law.

- **The triangle inequality FAILS for geometric triples** (Paper A, §5). Subadditivity holds on the parameter `u`, not on the geometry. You may not chain surplus bounds through an intermediate state as if `F` were a metric.
- **As deployed, the surplus functional is a LOWER BOUND** on the information carried by an ensemble, not the information itself. We log the floor, not the ceiling. Every number produced by this manual is a floor.
- **Paper A is pure mathematics and makes no empirical or metaphysical claim.**
- **Paper D:** surplus is not identified with Shannon entropy or mutual information in general.
- **Paper E: no thermodynamic negentropy law is proved.** Corollary 2.2 gives `F(u) + N_can(u) <= ln N`; higher static surplus does not imply higher canonical concentration.
- Applying the storage law to a room, a team, or a city is a **modelling decision**, badged as such every time it is made. The mathematics is proved. The mapping onto a social system is postulated, and Chapters 24 and 25 exist to make that postulate falsifiable.

The Engineering

1.5 Two recursions, one word

Consider two programs.

```
; PROGRAM A -- the stack
f(n):
    return 1 + f(n)          ; no base case, no decay
                             ; frames accumulate at rate 1 per call
                             ; terminates only by exhausting memory

; PROGRAM B -- the accumulator
Q <- 0
every period:
    Q <- (1 - delta)*Q + eta*S - C
                             ; state is ONE number, reused
                        ; each period discards a delta-fraction of the past
                             ; terminates by converging
```

Both are recursive: each step is defined in terms of the previous step. The word covers both. But the resource behaviour is opposite, and the reason is a single arithmetic fact.

In Program A, the contribution of the call made `t` steps ago still weighs 1. Nothing is ever discarded. The state is a *list* whose length grows without bound, and the machine's response to unbounded growth is to stop.

In Program B, the contribution of the input delivered `t` steps ago weighs `(1-delta)^t`. Read Theorem 3.1 again as a statement about memory:

```
Q_t = (1-delta)^t Q_0  +  sum_{k=0}^{t-1} (1-delta)^{t-1-k} (eta S_k - C_k)
      \_____________/     \_____________________________________________/
        the transient                  the driven accumulation
      (Book One Node 09)                (approaches Node 10)
```

Each historical input is present, but *discounted geometrically by its age*. The state is a single number because the geometric weighting makes the whole history summable. **Loss is what compresses an unbounded history into a bounded state.** That is the entire trick, and it is the same trick a capacitor performs and the same trick an exponential moving average performs.

Here is the weighting, computed. The entry is the share of the steady-state stock `Q_inf` that is attributable to inputs delivered within the last `m` periods, which is `1 - (1-delta)^m`:

periods back `m`	delta=0.02	delta=0.05	delta=0.10	delta=0.20
1	0.0200	0.0500	0.1000	0.2000
2	0.0396	0.0975	0.1900	0.3600
4	0.0776	0.1855	0.3439	0.5904
8	0.1492	0.3366	0.5695	0.8322
12	0.2153	0.4596	0.7176	0.9313
26	0.4086	0.7365	0.9354	0.9970
52	0.6503	0.9306	0.9958	1.0000

periods back `m`	delta=0.02	delta=0.05	delta=0.10	delta=0.20
104	0.8777	0.9952	1.0000	1.0000

Read the `delta = 0.10` column. Fifty-seven percent of what a system holds today was put there in the last eight periods; ninety-four percent in the last twenty-six. **A system with `delta = 0.10` is, functionally, a two-year institution no matter how old it is.** That is not a metaphor about culture. It is the third column of a table.

Now read the column that would exist at `delta = 0`. It is a column of zeros for every finite `m` except the trivial identity that all of history counts. Nothing is ever more recent than anything else. There is no such thing as *current practice* in a system with no forgetting, because the founding decision and this morning's decision carry identical weight forever.

1.6 Why the mirror tunnel ends

The folk image of infinite regress is two mirrors facing each other. Stand between them and you see a corridor of copies receding to a vanishing point, and the word people reach for is *infinite*.

It is not infinite, and the reason it is not infinite is `delta`.

A real mirror reflects a fraction `R < 1` of the light that hits it. After `n` bounces the intensity is `R^n`. The corridor terminates where `R^n` falls below the eye's threshold for distinguishing the image from the surround. Take that threshold at one percent of the source:

reflectivity `R`	bounces to 1%	bounces to 0.1%
0.70	12.9	19.4
0.80	20.6	31.0
0.90	43.7	65.6
0.95	89.8	134.7
0.98	228.0	341.9
0.99	458.2	687.3
0.999	4602.9	6904.3
1.000	infinite	infinite

(The arithmetic is `n = ln(threshold)/ln(R)`. This is the arithmetic of the illustration, not a measurement of any particular mirror; real thresholds depend on ambient light and on the observer.)

The corridor looks infinite and is finite, and the loss per bounce is what makes it finite. Set `R = 1` — a perfect mirror, no loss — and the sum of the image intensities genuinely diverges. The image that terrifies people about recursion is an image whose terror depends on assuming away the one physical parameter that every real mirror has.

`1 − delta` is the reflectivity of an institution. **The regress you fear is the behaviour of a system you have accidentally specified with perfect mirrors.**

1.7 What the contraction actually buys

Say `delta > 0` and hold the drive constant. Four distinct engineering properties follow, and it is worth separating them because organisations habitually want one and get all four, or want to suppress one and lose all four.

(a) A destination exists. `Q* = (eta S − C)/delta` is a number you can compute before you build. You are not iterating hopefully; you are iterating toward a value you can quote in the design review. If that value is unacceptable, you know before you start, and you change a parameter rather than trying harder.

(b) The destination does not depend on where you start. Uniqueness plus global convergence means the fixed point is a property of `(delta, eta, S, C)` alone. This is the mathematical form of a fact every practitioner knows: *you cannot fix a system's steady state by improving its launch.* A better kickoff changes `Q_0`, and `Q_0` decays out.

Here are three trajectories under identical parameters `delta = 0.10`, `eta S = 0.8243`, `C = 0.20`, so `Q* = 6.2433`:

t	from `Q_0 = 0`	from `Q_0 = 3`	from `Q_0 = 20`
0	0.000	3.000	20.000
1	0.624	3.324	18.624
2	1.186	3.616	17.386
5	2.557	4.328	14.366
10	4.066	5.112	11.040
20	5.484	5.849	7.916
30	5.979	6.106	6.826
40	6.151	6.195	6.447

The one that started at 20 is *falling*, and it is falling toward the same number the one that started at 0 is rising toward. The gap between any two trajectories is exactly `(1−delta)^t` times the initial gap: 20.00 at `t=0`, 6.97 at `t=10`, 2.43 at `t=20`, 0.848 at `t=30`, 0.0359 at `t=60`. **Two teams with identical parameters and wildly different histories become indistinguishable, and the rate at which they become indistinguishable is `delta`.**

(c) Errors decay. A perturbation is just a change in the state, and the state's memory of any change is `(1−delta)^t`. This is the property people are giving up when they try to build a system that never forgets:

periods since the error	delta=0	delta=0.01	delta=0.05	delta=0.10	delta=0.20
1	1.00000	0.99000	0.95000	0.90000	0.80000
5	1.00000	0.95099	0.77378	0.59049	0.32768
10	1.00000	0.90438	0.59874	0.34868	0.10737
25	1.00000	0.77782	0.27739	0.07179	0.00378
50	1.00000	0.60501	0.07694	0.00515	0.00001
100	1.00000	0.36603	0.00592	0.00003	0.00000

The first column is a column of ones. **A system with `delta = 0` is a system in which no mistake ever loses force.** Chapter 3 develops the consequence in full; state it here in its short form: a system with no forgetting cannot learn, because learning is the replacement of an old weighting by a new one, and replacement requires that the old weighting fall.

(d) The system is bounded, and the bound is knowable. Theorem 3.3 gives `Q_t <= max{Q_0, eta ln N / delta}`. For `eta = 0.5`, `N = 8`, `delta = 0.10`, that ceiling is `0.5 * 2.0794 / 0.10 = 10.397`. **No amount of interaction, however well-designed, pushes this system past 10.397 units of stored order.** That is the most useful sentence you can say to a sponsor who wants a tenfold improvement out of better meetings.

1.8 Growth without arrival: the `delta = 0` regime

Set `delta = 0` and keep everything else. The law becomes

```
Q_{t+1} = Q_t + eta S_t − C_t
```

with exact solution `Q_t = Q_0 + sum (eta S_k − C_k)` and uniform bound `Q_t <= Q_0 + t eta ln N` (Theorem 3.3, second case). Take the same drive as above, `eta S − C = 0.6243` per period:

t	Q_t with delta=0	Theorem 3.3 bound `Q_0 + t eta ln N`
0	0.000	0.000
1	0.624	1.040
5	3.122	5.199
10	6.243	10.397
20	12.487	20.794
50	31.216	51.986
100	62.433	103.972

Nothing here is *wrong*. The stock grows, monotonically, forever. But notice what has been lost:

- There is no fixed point, so there is no answer to "what will this settle at?"
- There is no attractor, so two systems that start differently stay different by exactly their initial gap, forever.
- There is no error decay, so every bad decision is permanent structure.
- The bound is linear in `t`, which means the only thing limiting the system is the calendar.

A system with `delta = 0` does not arrive. It only continues. **Unbounded accumulation is not success; it is the absence of a destination**, and it is the mathematical signature of an institution that has confused *having more rules* with *being better governed*.

Note carefully: this is *not* the failure people imagine when they say "infinite regress." It is worse in a subtler way. Nothing crashes. There is no stack overflow, no error message, no visible breakdown. The system just keeps piling up, and the ossification shows up much later, in Chapter 6's headroom equation, when `delta Q_t` finally eats the entire budget. Since `delta = 0`, that term never eats anything — instead the cost term `C_t` does, because every rule ever written must still be complied with, and compliance is `C`.

1.9 Rumination: recursion with eta near zero

Rumination is the most recognisable pathological loop in human experience, and it has an exact signature in the law: **interaction is happening, surplus is high, and almost none of it converts.**

Take a system that is genuinely thinking hard. N = 8, u* = 0.6, so S = ln(1 + 7 * 0.6) = ln 5.2 = 1.6487 — that is 79% of the theoretical maximum ln 8 = 2.0794. This is not a stupid system. It is generating real surplus every period.

Now set eta = 0.05: the interaction produces almost nothing durable. No artifact, no decision, no changed default, no trained person. Keep delta = 0.10 and C = 0.35 — thinking hard is expensive.

```
eta S = 0.05 * 1.6487 = 0.0824    vs    C = 0.35
```

Theorem 3.5 gives Q_inf = (0.0824 − 0.35)/0.10 = −2.676. **The equilibrium is negative, and Theorem 3.5 does not hide this: the regime is unsustainable or the linear model must be replaced.**

The trajectory from a healthy starting stock Q_0 = 6.2433:

t	0	1	2	3	5	8	12
Q_t	6.2433	5.3514	4.5487	3.8263	2.5909	1.1637	0 (floor)

Admissibility C_t <= (1−delta)Q_t + eta S_t requires 0.35 <= 0.9 Q_t + 0.0824, i.e. Q_t >= 0.2973. Below that, the model's own admissibility condition fails — which is the model telling you, correctly, that at this point something other than steady dissipation is happening.

The clinical picture matches the arithmetic exactly. Rumination is not a failure of effort; the effort is high. It is not a failure of insight; S is 79% of maximum. **It is a conversion failure: every period generates genuine surplus and then discards it, so the stock falls under its own maintenance cost.** The intervention that works is never "think harder" — that raises S, which is already high, and raises C, which is already fatal. The intervention that works is to raise eta: write one sentence down, make one decision, change one default, tell one other person. Chapter 4 is entirely about this parameter, and this is why it comes so early.

Note the diagnostic asymmetry that makes rumination so hard to see from inside: **a high-S, low-eta system feels productive and measures as declining.** The felt sense tracks S. The outcome tracks eta S. If you only ever ask people how the meeting went, you are measuring the wrong term.

1.10 Tradition: recursion with delta near zero and no input

The opposite pathology. Set delta at or near 0, set S = 0 — no live interaction, only transmission of what is already stored — and let C be small but nonzero, the ordinary cost of maintaining a practice nobody is adding to.

```
Q_{t+1} = Q_t − C_t,     C = 0.02,     Q_0 = 6.2433
```

t	0	10	25	50	100	200	312
Q_t	6.2433	6.0433	5.7433	5.2433	4.2433	2.2433	0.003

It reaches zero at `t = 312.2` periods. If a period is a year, that is a three-century institution that will be gone in three centuries and looks perfectly stable on any timescale a single custodian will observe.

This is the sealed vault. Its properties, read off the law:

- **It preserves superbly.** With `delta` near zero, `(1−delta)^t` is near 1 and the founding content survives intact. This is a real achievement and it is why traditions exist.
- **It cannot converge to anything new.** With `delta = 0` the map is not a contraction, so Banach gives nothing. There is no attracting state to move toward.
- **It cannot correct.** Column one of the error-persistence table: weight 1.000 forever. A founding error is load-bearing structure at year 300.
- **It decays only through cost, and only linearly.** Which means it fails *slowly and invisibly*, with no inflection point to alarm anyone.

Contrast the two pathologies precisely, because they are usually confused:

	rumination	tradition
S	high	zero
`eta`	near 0	irrelevant (nothing to convert)
`delta`	normal	near 0
felt experience	busy, exhausting, urgent	calm, solid, dignified
Q trajectory	falls fast, hits admissibility floor	falls slowly, linearly
the fix	raise `eta`	raise S (and accept some `delta`)
what it is good at	nothing	preservation

A system can be both at once, in different subsystems, and usually is: a ruminating leadership layer sitting on top of a sealed operational layer. The survey in Chapter 24 is designed to catch that.

1.11 Feedback is not regress

The distinction that resolves the confusion in ordinary language is between **feedback** and **regress**.

```
REGRESS                              FEEDBACK
-------                              --------
each step ADDS a new frame           each step UPDATES one state
history stored explicitly            history stored as a discounted sum
weight of step k = 1 forever         weight of step k = (1−delta)^age
resource use grows with t            resource use constant in t
terminates by exhaustion             terminates by convergence
no fixed point                       unique fixed point (delta > 0)
```

"we keep going in circles" "we are converging on a practice"

Here is the diagram of the circuit, which will recur in every part of this manual:

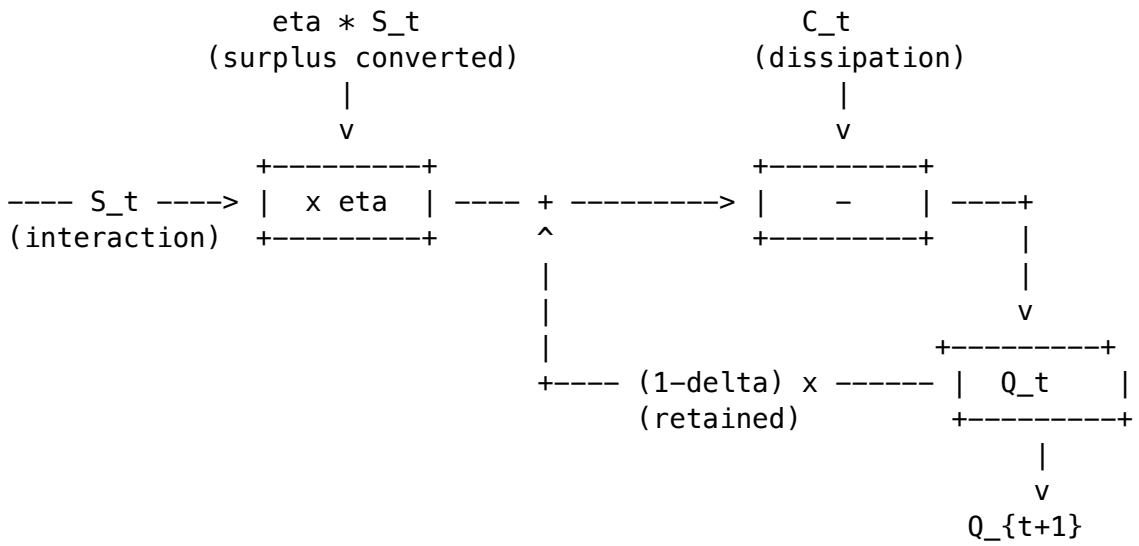

> The loop from Q back into the summing junction through the gain (1−delta) IS the recursion. delta is the loss in that path. If delta = 0 the loop gain is exactly 1 and the integrator never settles.

A control engineer reading that diagram will say: loop gain 1, marginally stable, pole on the unit circle. That is precisely correct, and it is the same statement as "no Banach contraction, no fixed point." **The organisational complaint "we go round in circles" and the control-theoretic complaint "the pole is on the unit circle" are the same complaint.**

The fix in control engineering is not to break the loop. It is to move the pole inside the unit circle — that is, to introduce loss. This manual's whole project is the organisational version of that move.

1.12 The horizon of arrival

Corollary 3.2 says growth continues exactly while `eta S_t > delta Q_t + C_t`. With the drive held constant this is a statement about a threshold in `Q`:

```
grow while  Q_t < (eta S − C)/delta = Q_inf
```

which is to say the fixed point is not only where you end up, it is the boundary between the growing regime and the shrinking regime. Below it you grow, above it you shrink, and the "above it" case is not hypothetical — it is what happens to an organisation whose drive falls (a key person leaves, S drops) while its stock stays high.

Worked, with `delta = 0.10`, `C = 0.20`, `eta S = 0.8243`:

Q_t	delta Q_t + C_t	vs eta S_t = 0.8243	verdict
0.0000	0.2000	less	grow
2.0000	0.4000	less	grow
4.0000	0.6000	less	grow
6.0000	0.8000	less	grow

Q_t	delta Q_t + C_t	vs eta S_t = 0.8243	verdict
6.2433	0.8243	equal	**hold — this is Q_inf**
7.0000	0.9000	greater	shrink
10.0000	1.2000	greater	shrink

Two things to take from this table, both structural.

First, **the maintenance term delta Q_t grows with success.** At Q = 0 the system spends 0.20 per period on upkeep. At Q = 6.24 it spends 0.824. Same organisation, same delta, four times the upkeep, purely because it has more to keep. This is the seed of Chapter 6's central result: success consumes its own headroom.

Second, the headroom gauge along the same trajectory, h_t = eta ln N − delta Q_t − C_t with eta ln N = 0.5 ∗ 2.0794 = 1.0397:

Q_t	0	2	4	6	6.2433	8	10
h_t	0.8397	0.6397	0.4397	0.2397	0.2154	0.0397	**-0.1603**

h hits zero at Q = 8.397. Beyond that point, **Corollary 3.4 applies in its hard form: no arrangement of people, no meeting design, no facilitation technique, no interaction geometry whatsoever prevents this system's stored order from decreasing.** You are not looking at a motivation problem. You are looking at an inequality.

1.13 Hyperrecursive symmetry

POSTULATE, badged. The storage law is used in this manual at every scale — a desk, a room, a team, a building, an institution, a city — with only (delta, eta, N, C) changing between them. This is a claim about **model form**, not about physics. It asserts that the same difference equation is a useful description at each scale. It does not assert that a city is a large room, and it does not assert any conservation law, any thermodynamic principle, or any mechanism connecting the scales.

The claim earns its place by being falsifiable at each scale independently. Chapter 25 gives the protocol: fit Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t on training data with 0 <= a <= 1, b >= 0, freeze the model, and report held-out performance against four baselines. If the fitted a is not in [0,1], or the model does not beat a matched autoregression and a shuffled-input control, the form is wrong *at that scale*, and you say so.

What the symmetry buys, when it holds, is transfer of method rather than transfer of numbers. The delta of a whiteboard and the delta of a constitution differ by orders of magnitude. The *procedure* for measuring both is identical, and that is Chapter 24.

1.14 The method of this manual

Every chapter after this one is an instrument for one of five quantities, or an application of the law to one class of system. The method is uniform:

Step 1 — Identify `Q` and choose a proxy `Y`. What is the stored order you care about? What observable stands in for it? **Justify that `Y` is not merely a repackaging of the input `U`** (Paper E, §6.3). This is the most common measurement error in the field and Chapter 24 spends several pages on it.

Step 2 — Measure `delta`. Find or create a window of near-zero input and fit the decay. `delta = 1 - (Y_k/Y_0)^{1/k}`, or better, a log-linear regression on several points. Chapter 3.

Step 3 — Measure `eta`. Attributable change in `Y` per unit of surplus delivered. `eta = delta_Y_attributable / S_input`. Chapter 4.

Step 4 — Enumerate `N` and estimate `u`. How many genuinely distinct categories, and how non-parallel are the states within them? Chapters 7 through 11.

Step 5 — Tally `C`. All of it, including the parts nobody bills for: coordination, commute, compliance, emotional labour, context switching. Chapter 5.

Step 6 — Compute `h_t = eta ln N - delta Q_t - C_t` and read the regime. Chapter 6.

Step 7 — Choose the cheapest lever that moves `h` in the direction you need. Chapter 26 ranks them: `eta` is usually cheapest, `C` next, `u` next, `N` most expensive and least rewarding because its contribution is logarithmic.

Step 8 — Preregister the prediction and the falsification criteria before you intervene. Chapter 25. **A framework that cannot fail cannot be trusted, and this one names exactly how it fails.**

Each chapter carries the same five-part armature: **The Law** (proved or postulated, badged), **The Engineering** (the body), **The Specification** (numbered, checkable), **The Failure Modes** (named faults with diagnostic signatures and fixes), and **Worked Example** (one case carried end to end with real arithmetic). Chapter 27 collects every named failure mode into a lookup catalogue and Appendix D indexes it by symptom.

One rule governs all of it, from CANON §5 and repeated here because it is the constraint that makes the rest of the manual usable: **any intervention that removes a person's ability to decline is out of scope.** Not discouraged. Out of scope. The method in this book works by lowering the cost of a better arrangement until people move to it voluntarily; its worst failure mode is that nobody moves. That is a design property, not a limitation, and Chapter 20 makes it explicit.

The Specification

Requirements for any system this manual is applied to. Each is checkable.

S1.1 — Loss must be positive and named. Every modelled subsystem shall have `delta > 0`, and the mechanism of loss shall be identified by name (turnover, forgetting, wear, drift, deprecation). A subsystem for which no loss mechanism can be named has not been modelled; it has been assumed.

S1.2 — The fixed point shall be computed before the system is built. For the intended operating parameters, compute `Q_inf = (eta F(u*) - c*)/delta` and record it in the design record. If `eta F(u*) < c*`, the design is rejected at this line and no further work proceeds until a parameter changes.

S1.3 — The uniform bound shall be quoted. Record `eta ln N / delta` as the ceiling on stored order (Theorem 3.3). Any commitment that exceeds it is void.

S1.4 — Convergence time shall be stated, not assumed. Record the number of periods to reach 95% and 99% of `Q_inf` from the actual starting stock. Formulae and tables are in Chapter 2. A programme whose sponsor's patience is shorter than its `t_95` is misdesigned regardless of its fixed point.

S1.5 — Admissibility shall be checked each period. Verify `C_t <= (1−delta) Q_t + eta S_t`. A violation is escalated as a model-breakdown event, not clipped silently.

S1.6 — Headroom shall be measured before any increment is promised. Compute `h_t = eta ln N − delta Q_t − C_t`. Any requested increment `r_t` shall satisfy `r_t <= h_t` (Corollary 4.2). Promises exceeding `h_t` are not ambitious; they are arithmetically unavailable.

S1.7 — The order proxy shall be justified as distinct from the input. Written justification that `Y` is not a repackaging of `U` (Paper E, §6.3) shall exist before any fitting is performed.

S1.8 — Postulates shall be badged in the sentence that states them. Any application of the storage law to a social or physical system is a modelling decision and shall be marked as such at the point of use, not in a footnote.

S1.9 — Every number shall be computed or cited. No figure appears in a design record without either the arithmetic that produced it or the source that reported it. Where a figure would require research not yet done, the record states what would need measuring instead.

S1.10 — Reversibility. Any intervention shall have a defined stop condition and a defined return path. A change that cannot be undone is a change to `delta` as well as to whatever it was aimed at, because it becomes permanent structure.

S1.11 — Declining shall remain available. No specification in this manual is satisfied by a mechanism that removes a participant's ability to opt out. (CANON §5.4.)

The Failure Modes

Each entry: **NAME** — clinical description; damaged parameter; diagnostic signature; measurement that detects it; fix. Chapter 27 holds the full catalogue; these are the ones this chapter is responsible for.

THE UNTERMINATED REGRESS — the system reruns the same deliberation without converting any of it, and each pass restates the previous pass rather than updating a single stored position.

- *Parameter damaged:* `eta` (to zero) and effectively `delta = 0` on the deliberation record.
- *Diagnostic signature:* meeting `n+1` opens by re-establishing what meeting `n` established. Agenda items reappear with the same wording. Participants can recall the discussion but cannot point to an artifact.
- *Measurement:* count decisions recorded per hour of deliberation. Count agenda-item recurrences across consecutive meetings. Both near zero and high, respectively.

- *Fix:* force a single mutable stored state — one document, one decision log, one register — that each pass must update rather than restate. This converts the stack into an accumulator by construction. See Chapter 4 and Chapter 21 (put the state in the medium).

THE RUMINATION LOOP — high-surplus interaction with near-zero conversion and nonzero cost; stock falls while the system feels maximally engaged.

- *Parameter damaged:* `eta`.
- *Diagnostic signature:* `S` is high (real disagreement, real analysis, real expertise in the room) and `Q` is falling. Participants report the sessions as valuable. Nothing downstream changes.
- *Measurement:* `eta = delta_Y_attributable / S_input` over several periods. Compare `eta S` against `C` directly. If `eta S < C`, Theorem 3.5 gives a negative equilibrium and you have your answer.
- *Fix:* raise `eta`, never `S`. Smallest effective moves: a written decision at the end of each session; a named owner; a changed default; a trained second person. Chapter 4.

THE SEALED VAULT — `delta` near zero with no live input; excellent preservation, zero adaptation, slow linear decline through cost alone.

- *Parameter damaged:* `S` (to zero); `delta` suppressed below useful.
- *Diagnostic signature:* the system's answer to a novel question is a citation to an old decision. Procedures whose original rationale nobody can state are still enforced. Founding errors are still load-bearing.
- *Measurement:* age distribution of the active rule set. Fraction of stored order whose rationale can be reconstructed by a current member. Rate of rule retirement (usually exactly zero).
- *Fix:* introduce input (`S > 0`) and introduce *selective* loss. Sunset clauses, review cycles, and scheduled retirement are `delta` restored deliberately on the parts that should decay. Chapters 3 and 23.

THE UNBOUNDED LEDGER — `delta = 0` with continuing input; stock grows linearly and forever, with no fixed point and no error decay. Every rule ever written is still in force.

- *Parameter damaged:* `delta` (to zero), with `C` rising as a consequence.
- *Diagnostic signature:* rule count, policy count, or documentation volume rises monotonically and has never fallen. Compliance cost rises with no corresponding rise in capability. Nobody can state the current state of practice without reading everything.
- *Measurement:* plot cumulative rule count against time; look for any decrease. Measure onboarding time to competence — it grows with `Q` when `delta = 0`.
- *Fix:* institute retirement. The `delta = 0` system has no attractor; giving it one requires removing structure on purpose. Chapter 23 covers deliberate shedding.

THE FORGOTTEN COST — the model is specified with `S`, `eta`, and `delta`, and `C` is left out because nobody owns it. The design predicts growth and the system declines.

- *Parameter damaged:* `C` (unmeasured, therefore unbudgeted).

- *Diagnostic signature:* predicted `Q_inf` far exceeds observed steady state. Predicted growth, observed flatness.
- *Measurement:* tally every recurring draw on the system, including unbilled ones — coordination, context switching, commute, compliance, emotional labour. Chapter 5 gives the enumeration.
- *Fix:* re-fit with `C` included and re-check `eta F(u*) > c*`. Frequently the correct action is to cut `C` rather than to add `S`, because `C` enters linearly and `S` enters through a logarithm.

THE UNBADGED POSTULATE — a modelling assumption is presented with the authority of a theorem.

- *Parameter damaged:* none directly; it damages the warrant of every number downstream.
- *Diagnostic signature:* a claim of the form "the law shows that this room will…" where the law shows nothing of the kind without an unstated mapping from room to model.
- *Measurement:* audit every quantitative claim in the design record against CANON §1. Anything not traceable to a theorem is a postulate and must carry its badge.
- *Fix:* restate with the badge in the same sentence. **The mysticism is permitted; it is permitted in the open, with its badge on.**

Worked Example

A twelve-person research group

The system. A twelve-person applied research group inside a larger organisation. Period: one week. The question is whether a proposed change to the weekly seminar can deliver the improvement the director has promised upward.

Step 1 — `Q` and its proxy. `Q` is the group's stored methodological order: techniques that any member can execute correctly without the originator present. The proxy `Y` is the count of techniques that pass a monthly spot-check, where the check is performed by a member who did not develop the technique. This proxy is not a repackaging of the input `U`, because `U` describes the *disagreement geometry* across research categories in a given week and `Y` describes *transferable executable capability* measured a month later by a different person (Paper E, §6.3). That justification is written into the design record before any fitting, per S1.7.

Step 2 — `N` and `u*`. The group spans eight genuinely distinct methodological categories, so `N = 8` and the per-interaction ceiling is `ln 8 = 2.0794`. Disagreement audits and error- decorrelation measurements across category pairs (Chapter 9's protocol) give a working estimate `u* = 0.60`.

```
N_eff = 1 + (N−1) u* = 1 + 7 * 0.60 = 5.20
S = F(u*) = ln(5.20) = 1.6487
```

The roster says twelve people in eight categories. The audited effective count is 5.20. The group is buying eight categories' worth of salary and receiving 5.20 categories' worth of non-redundancy. Here is the full sensitivity, which the director should see:

u*	N_eff = 1 + 7u*	S = ln N_eff	% of ceiling ln 8
0.00	1.00	0.0000	0%
0.10	1.70	0.5306	26%
0.25	2.75	1.0116	49%
0.50	4.50	1.5041	72%
0.60	**5.20**	**1.6487**	**79%**
0.75	6.25	1.8326	88%
0.90	7.30	1.9879	96%
1.00	8.00	2.0794	100%

Note the concavity (Paper A, 3.3) doing its work: moving u* from 0.00 to 0.25 buys 1.0116 nats; moving it from 0.75 to 1.00 buys 0.2468. **The first quarter of non-parallelism is worth four times the last quarter.** Chapter 7 develops this as the merciful clause.

Step 3 — `eta`. Measured over eleven weeks as attributable change in Y per unit of S delivered: `eta = 0.50`. So `eta ln N = 0.5 * 2.0794 = 1.0397` is the absolute per-week input ceiling, and the achieved input is `eta S = 0.5 * 1.6487 = 0.8243`.

Step 4 — `delta`. A four-week shutdown gave a clean zero-input window; log-linear fit on the decay gives `delta = 0.10` per week. Half-life `ln 2 / (-ln 0.9) = 6.58` weeks. Chapter 3 gives the procedure and Chapter 3's worked example gives a fuller version of this fit.

Step 5 — `C`. Tallied at `C = 0.20` per week: seminar preparation, the coordination overhead of scheduling across eight categories, and the write-up burden.

Step 6 — the fixed point. By Theorem 3.5,

```
Q_inf = (eta F(u*) - c*) / delta = (0.8243 - 0.20) / 0.10 = 6.2433
```

and by Theorem 3.3 the uniform ceiling is `eta ln N / delta = 1.0397 / 0.10 = 10.397`. **The group converges to 6.2433 and cannot under any interaction geometry exceed 10.397.**

Step 7 — convergence. Starting from `Q_0 = 0`:

week	Q_t	% of Q_inf
0	0.0000	0.00%
1	0.6243	10.00%
2	1.1862	19.00%
4	2.1471	34.39%
6	2.9254	46.86%
8	3.5558	56.95%
10	4.0664	65.13%
12	4.4800	71.76%
15	4.9579	79.41%
20	5.4843	87.84%
25	5.7951	92.82%
29	5.9492	95.29%
30	5.9786	95.76%

Ninety-five percent of the destination is reached in week 29; the exact figure from `t_95 = ln(0.05)/ln(0.90)` is 28.43 weeks. Ninety-nine percent takes `ln(0.01)/ln(0.90) = 43.71` weeks. **The director's promised review is at week 12, when the system will be at 71.76% of its own steady state and still climbing.** That is not underperformance. It is the step response of a first-order system evaluated before it settles, and Chapter 2 exists so that nobody makes this mistake in a design review.

Step 8 — headroom, and what may be promised. With `eta ln N = 1.0397`:

`Q_t`	`h_t = 1.0397 − 0.10 Q_t − 0.20`	regime
0.0000	0.8397	growth
2.0000	0.6397	growth
4.0000	0.4397	growth
6.2433	0.2154	growth (at `Q_inf`, still positive)
8.0000	0.0397	growth, nearly exhausted
8.3972	0.0000	**maintenance only**
10.0000	-0.1603	**unsustainable — no geometry saves it**

The maximum single-week increment the group can be promised at its current stock of 4.48 is `r_max = 1.0397 − 0.448 − 0.20 = 0.3917` (Corollary 4.2). By Proposition 4.1, achieving an increment `r_t` requires `S_t >= (delta Q_t + C_t + r_t)/eta`, feasible only if `delta Q_t + C_t + r_t <= eta ln N`. At `Q_t = 4.48` and `r_t = 0.50`, the requirement is `S_t >= (0.448 + 0.20 + 0.50)/0.50 = 2.296`, which exceeds `ln 8 = 2.0794`. **The requested increment is infeasible, and the proof is one line of arithmetic.** The correct report upward is not "we will try"; it is "0.50 per week is above `r_max = 0.3917`, here is the calculation, and here are the three parameters that would have to move."

Step 9 — the counterfactuals, computed.

If `delta` were 0. The group would grow linearly at `0.6243` per week without limit — 6.243 at week 10, 12.487 at week 20, 62.433 at week 100 — with no fixed point, no error decay, and no answer to "what will this settle at?" It would also never lose a single obsolete technique, and its spot-check burden would grow without bound, which shows up as a rising `C` in year three. The group's methodological register would become an archive that must be read in full before anyone can act. **The absence of loss does not produce a better group; it produces a group that cannot say what it currently does.**

If `eta` fell to 0.05 — the seminar continues, the arguments are as good as ever, and nobody writes anything down. Then `eta S = 0.0824 < C = 0.20`, `Q_inf = −2.676`, and the stock falls from 6.2433 to below the admissibility floor within twelve weeks. The seminar would feel *better* during this period, because `S` is unchanged and the discussions are as rich as they ever were.

If `S` fell to 0 and `delta` were driven to 0 — the sealed vault. `Q` falls linearly at `C = 0.20` per week and reaches zero at week 31.2, with the group defending techniques whose justification nobody present can reconstruct.

Step 10 — the intervention, and its falsification criterion. The binding constraint is not $u*$, which at 0.60 already delivers 79% of the ceiling and whose remaining headroom is worth only 0.43 nats at enormous recruiting cost. It is not `N`, because raising `N` from 8 to 16 buys exactly `ln`

2 = 0.6931 extra ceiling for a doubling of genuinely distinct methodological categories — see Chapter 8. The binding constraint is the pair (eta, C), and of the two, eta is cheaper to move (Chapter 26's lever ranking).

The intervention: every seminar closes with one written, owned, dated decision, and one named second person who can execute the technique discussed. Predicted effect: eta from 0.50 to 0.60, C unchanged. Then eta S = 0.9892, Q_inf = (0.9892 − 0.20)/0.10 = 7.892, a 26.4% increase in steady-state stored order, reached to 95% in another 28.4 weeks from the point of change.

Preregistered falsification, per Chapter 25: fit Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t on weeks 1-30, freeze, and predict weeks 31-60. The intervention is falsified if the fitted a does not lie in [0,1], if b does not increase, or if the frozen model fails to beat all four required baselines on held-out data — persistence, direct-input, matched autoregressive, and the shuffled-input control. **The shuffled-input control is the one most often omitted and the one most often decisive**, because it is the only baseline that tests whether the *timing* of interaction matters or merely its volume.

What this chapter has established

The equation this manual runs on is a recursion. It does not overflow, it does not regress, and it does not need a base case, because it has something better: a loss term that discounts the past geometrically and thereby compresses an unbounded history into a single bounded number.

That loss term makes the update map a contraction. Banach then supplies existence, uniqueness, and global convergence, which is to say: a destination, a destination independent of history, and arrival at it from anywhere. Remove the loss and all three disappear simultaneously — not because anything crashes, but because there is nothing left to converge to.

The leak is not the flaw. The leak is what makes arrival possible.

Chapter 2 develops the circuit identity in full — pole, time constant, settling time, DC gain, frequency response — and shows why a ritual on a fixed cadence outperforms an intense one-off by an amount that can be computed in advance. Chapter 3 takes delta itself as the object of engineering: what it is made of, how to measure it, how to drive it down where you want retention, and how to drive it *up* where you want a system that can still change its mind.

Chapter 2 . The Leaky Integrator

Chapter 1 established that the storage law is a recursion with loss and that loss is what makes arrival possible. This chapter cashes that in.

The claim is stronger than an analogy. **The storage law is not *like* a first-order low-pass filter; in discrete time it *is* one**, term for term, with the same pole, the same time constant, the same step response, the same settling time and the same DC gain. Every result in the standard theory of first-order discrete systems is therefore available without modification, and none of it has to be rediscovered in sociological language.

That matters practically, because the standard theory answers questions organisations ask constantly and answer badly:

- *When should we review this initiative?* Settling time.
- *Will the offsite work?* Impulse response, compared against the input ceiling.
- *Is a big push better than a steady cadence?* Frequency response plus a saturation bound.
- *Why did the change take a year to show up?* The pole.
- *Why is a two-year-old programme still climbing?* `t_95` at the measured `delta`.

This chapter derives the correspondence, computes the tables, and gives the design rules that come out of them. The load-bearing result at the end is this: **ritual is a periodic forcing function on a low-pass filter, and that — not enthusiasm, not commitment — is why cadence beats intensity.**

The Law

2.1 The identity

The storage law (Paper E, §3):

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,    S_t = F(U_t) = ln(1 + (N−1) U_t)
```

Set `x_t := eta S_t − C_t` for the net drive. Then

```
Q_{t+1} = a Q_t + x_t,      a := 1 − delta
```

which is the canonical first-order discrete-time system. Its transfer function, taking the z transform with `Q_0 = 0`, is

```
H(z) = Q(z) / X(z) = 1 / (z − a) = z^{−1} / (1 − a z^{−1})
```

with a single **pole at `z = a = 1 − delta`**. The system is bounded-input bounded-output stable if and only if `|a| < 1`, that is, if and only if `0 < delta < 2`. Since `delta` lies in `[0,1]` by construction, the binding case is `delta > 0`.

This is an algebraic identity, not a modelling assumption. The badge belongs one level up: the decision to describe a room, a team, or an institution by this difference equation *is* a modelling postulate (CANON §1.4; badged wherever used). Given the equation, the circuit facts follow with no further assumption.

2.2 The correspondence table

storage law	circuit	control theory
`Q_t`	charge on the capacitor, `q = C_cap V`	state
`delta`	leakage conductance across the capacitor	`1 −` pole
`1 − delta`	retention per period	pole location a
`eta S_t`	current in from the source	driven input
`C_t`	current drawn by the load	disturbance
`eta`	source transconductance	input gain
`eta / delta`	DC gain of the RC section	DC gain `H(1)`
`Q_inf`	steady-state charge	DC steady state
`(1−delta)^t Q_0`	initial charge bleeding off	transient / natural response
`1 / delta` (periods)	`tau = R C_cap` (seconds)	time constant
`S_t <= ln N`	source compliance limit	input saturation
`C_t <= (1−delta)Q_t + eta S_t`	capacitor cannot supply more than it holds	admissibility

2.3 Continuous RC, and where `delta` comes from

The continuous first-order low-pass is

```
tau dV/dt = V_in(t) − V(t),      tau = R C_cap
```

Euler-discretise with step `dt`:

```
V(t+dt) = V(t) + (dt/tau) (V_in − V(t)) = (1 − dt/tau) V(t) + (dt/tau) V_in
```

Matching coefficients against `Q_{t+1} = (1−delta)Q_t + ...` gives

```
delta = dt / tau      i.e.      tau = dt / delta = 1/delta   periods
```

`delta` is not a property of the system alone. It is the ratio of your observation period to the system's own time constant. A `delta` of 0.05 measured weekly and a `delta` of 0.19 measured monthly can be the same system: `1 − (1−0.05)^4 = 0.1855`. §2.11 gives the conversion rules and the errors that come of ignoring them.

2.4 DC gain and the steady state

Evaluate the transfer function at `z = 1` (zero frequency, i.e. a constant input):

```
H(1) = 1/(1 − a) = 1/delta
```

so a constant net drive x produces steady state `Q_inf = x/delta`. Substituting `x = eta F(u∗) − c∗` recovers Paper E's Theorem 3.5 exactly:

```
Q_inf = (eta F(u∗) − c∗) / delta
```

Theorem 3.5 is the DC gain of the circuit. If you measure surplus in nats and cost in the same units per period, the DC gain from surplus to stored order is `eta/delta`: each sustained unit of surplus per period buys `eta/delta` units of permanent stock.

Two consequences worth stating separately because they are constantly conflated:

- `eta` sets how much of each period's interaction becomes structure.
- `eta/delta` sets how much stock a *sustained* rate of interaction supports.

A system with `eta = 0.4` and `delta = 0.08` has a DC gain of 5.0. Halving `delta` to 0.04 doubles the DC gain to 10.0 without touching `eta` at all. **Retention is a gain term, and organisations almost never treat it as one.**

A note on deliberately raising `C`, sharpened by Chapters 41 and 42. The DC gain makes cost look like a term you would never willingly increase, and the popular instruction *get out of your comfort zone* looks, from here, like straightforward self-harm. That reading is too coarse, and Book One's flat rejection of the slogan is corrected here rather than repeated. The measured answer splits in two:

- **A cost rise applied globally is strictly harmful, monotonically.** It lowers `eta F(u*) − c*`, lowers `Q_inf`, and lengthens every departure. Chapter 41 §41.7 computes an operator's wait to fund a departure rising from **32.0 weeks at baseline to 42.5, then 89.2, then never** — because above a critical operating cost `C_crit = eta M − lambda E_b` (0.082 Q-units/week in that case) the departure is unreachable at any level of desire and unlimited time.
- **A cost rise confined to the origin genuinely shortens the wait**, because the barrier is measured as excess cost over the starting configuration: raising the origin lowers the budget and the barrier together. Chapter 41 §41.7.1 measures **32.0 weeks down to 9.9** — bought at a permanent fund **37% smaller** (`Λ_inf` falling from 4.32 to 2.72).

The question is therefore narrower than a denunciation: is it a cost you can leave behind, or a cost you carry? Local discomfort at the origin works. Global discomfort — attention, sleep, health, regulation, carry — follows you out of the basin and destroys the platform. See Chapter 41 §41.7.1 and Chapter 42's Worked Example, which runs the comparison on a measured field.

2.5 Honesty clauses

- The circuit identity is exact for the *equation*. Whether the equation describes your team is a postulate, badged, and falsifiable by the protocol in Chapter 25.
- As deployed the surplus functional is a **lower bound** on the ensemble's information (Paper A/D), so every `Q_inf` computed here is a floor, not a forecast.
- The model is **linear**. Real systems saturate, have thresholds, and have state-dependent `eta`. Theorem 3.5 flags its own breakdown: if `eta F(u*) < c*` the equilibrium is negative and either the regime is unsustainable or the linear model must be replaced.
- **No thermodynamic claim is made.** The capacitor is a mathematical correspondence, not a physical mechanism, and Paper E proves no negentropy law.

The Engineering

2.6 The pole, and what it feels like from inside

The pole `a = 1 − delta` is the multiplier applied to the whole stored state each period. Three regimes, and all three occur in real organisations:

```
|a| < 1   (delta > 0)      STABLE. Contraction. Unique fixed point.
                           Transients decay as a^t. Errors forgotten.
                          --> Chapter 1: this is why iteration arrives.

 a  = 1   (delta = 0)      MARGINALLY STABLE. Pure integrator.
                        No fixed point. Any constant input ramps forever.
                           Initial condition never decays.
                       --> "we go round in circles"; the unbounded ledger.

|a| > 1  (delta < 0)     UNSTABLE. Not admissible here: delta lies in [0,1].
                      Would correspond to stored order that amplifies itself
                           without input. The law does not permit it.
```

The third case is worth a sentence even though the law excludes it, because organisations *describe* themselves as being in it — "the culture reinforces itself," "it snowballs." Under the storage law that behaviour is not `a > 1`; it is a high `eta` combined with an `S` that rises with `Q`, which is a *nonlinearity in the drive*, not instability in the pole. If your system genuinely shows self-amplification, the linear model is the wrong model and Chapter 25's fitting procedure will report `a > 1`, which is your signal to stop and re-specify rather than to celebrate.

2.7 The time constant

Two definitions are in circulation and they differ by a few percent at ordinary `delta`:

```
tau_approx = 1 / delta                (the engineering rule of thumb)
tau_exact  = -1 / ln(1 - delta)       (the true e-folding time of a^t)
```

`tau_exact` is the number of periods for the transient `a^t` to fall to `1/e`. `tau_approx` is its first-order approximation and is the number everyone uses. They agree to better than 1% for `delta <= 0.02` and to about 5% at `delta = 0.10`. Below `delta = 0.30` either is fine for design work; above it, use the exact form, and above `delta = 0.5` stop using time constants at all and read the settling times directly.

Half-life follows immediately:

```
t_half = ln 2 / (-ln(1 - delta)) = 0.6931 * tau_exact
```

2.8 The master table: `delta`, pole, time constant, settling

This is the table to pin above the desk. All figures computed from `t_p = ln(1-p) / ln(1-delta)`.

delta	pole 1−delta	tau = 1/delta	tau_exact	half-life	t_63	t_95	t_99	3 tau	5 tau
0.001	0.999	1000.00	999.50	692.80	999.17	2994.23	4602.87	3000.0	5000.0
0.005	0.995	200.00	199.50	138.28	199.43	597.65	918.73	600.0	1000.0
0.01	0.990	100.00	99.50	68.97	99.47	298.07	458.21	300.0	500.0
0.02	0.980	50.00	49.50	34.31	49.48	148.28	227.95	150.0	250.0
0.05	0.950	20.00	19.50	13.51	19.49	58.40	89.78	60.0	100.0

delta	pole 1− delta	tau = 1/delta	tau_exact	half-life	t_63	t_95	t_99	3 tau	5 tau
0.08	0.920	12.50	11.99	8.31	11.99	35.93	55.23	37.5	62.5
0.10	0.900	10.00	9.49	6.58	9.49	28.43	43.71	30.0	50.0
0.15	0.850	6.67	6.15	4.27	6.15	18.43	28.34	20.0	33.3
0.20	0.800	5.00	4.48	3.11	4.48	13.43	20.64	15.0	25.0
0.30	0.700	3.33	2.80	1.94	2.80	8.40	12.91	10.0	16.7
0.40	0.600	2.50	1.96	1.36	1.96	5.86	9.02	7.5	12.5
0.50	0.500	2.00	1.44	1.00	1.44	4.32	6.64	6.0	10.0
0.60	0.400	1.67	1.09	0.76	1.09	3.27	5.03	5.0	8.3
0.70	0.300	1.43	0.83	0.58	0.83	2.49	3.82	4.3	7.1
0.80	0.200	1.25	0.62	0.43	0.62	1.86	2.86	3.8	6.3
0.90	0.100	1.11	0.43	0.30	0.43	1.30	2.00	3.3	5.6

The rule of thumb `3 tau` to 95% and `5 tau` to 99% is conservative in discrete time, which is the safe direction. Checking it exactly, the fraction of the step reached at `k tau` (with `tau = 1/delta`) is:

delta	at 1 tau	at 2 tau	at 3 tau	at 4 tau	at 5 tau
0.02	0.6358	0.8674	0.9517	0.9824	0.9936
0.05	0.6415	0.8715	0.9539	0.9835	0.9941
0.10	0.6513	0.8784	0.9576	0.9852	0.9948
0.20	0.6723	0.8926	0.9648	0.9885	0.9962
continuous 1 − e^{−k}	0.6321	0.8647	0.9502	0.9817	0.9933

So `3 tau` gets you 95.2% to 96.5% and `5 tau` gets you 99.4% to 99.6%, always slightly ahead of the continuous-time reference. Use `3 tau` for "substantially arrived" and `5 tau` for "settled."

Design rule R2.1: no initiative is evaluated before `3 tau` periods have elapsed. Evaluating a first-order system at `1 tau` and finding it 65% of the way there is not a finding about the initiative. It is a finding about the reviewer's calendar.

2.9 Step response: a policy change

A **step** is a permanent change in the drive: a policy that starts and stays, a new standing meeting, a hire, a rule. Formally `x_t = x` for all `t >= 0`, `x_t = 0` before.

```
Q_t = Q_inf (1 − a^t) + Q_0 a^t,     Q_inf = x/delta
```

Two exponentials with the same rate: the old state bleeding out, the new state filling in. There is no overshoot, no oscillation, no ringing. **A first-order system with a real pole in `(0,1)` approaches its new steady state monotonically from whichever side it starts on.**

Worked, `delta = 0.08`, `x = 0.3511` per week, so `Q_inf = 4.3888`:

week	Q_t	% of Q_inf
0	0.0000	0.00%
1	0.3511	8.00%
2	0.6741	15.36%
4	1.2447	28.36%
6	1.7276	39.36%
8	2.1364	48.68%
12	2.7752	63.23%
18	3.4104	77.71%
24	3.7955	86.48%
26	3.8867	88.56%
36	4.1707	95.03%
48	4.3086	98.17%
52	4.3314	98.69%
60	4.3593	99.33%

At the six-month review (week 26) the system is at 88.6% of where it is going. At the annual review (week 52) it is at 98.7%. **The visible "slowdown" between month six and month twelve is not the initiative losing momentum; it is `a^t` doing exactly what `a^t` does.** An organisation that reads that curve as fatigue and responds by adding input has misdiagnosed a pole as a morale problem.

2.10 Impulse response, and the ceiling that kills the offsite

An **impulse** is a single period of input followed by nothing: the offsite, the launch event, the one-time grant, the visiting expert.

```
Q_t = A a^{t}   for t >= 0, from a single input A at t = 0
```

Pure decay. The whole contribution of an impulse is `A/(1)` at the moment it lands and `A * a^t` thereafter, and its total contribution to the running integral over all time is `A/delta`. That last figure is what makes intensity *sound* competitive: an impulse and a sustained drive with the same total input have the same integrated effect over infinite time. In a linear time-invariant system this is exactly true, and honesty requires saying so plainly.

The reason intensity nevertheless loses is not the linear algebra. It is the saturation bound.

Paper D's Corollary 2.2 gives `0 <= F(u) <= ln N`, so the input in any single period is bounded:

```
eta S_t <= eta ln N
```

You cannot deliver a year's worth of surplus in a weekend, because the per-period ceiling is a hard constraint of the law, not a matter of stamina. Some ceilings, computed:

eta	N	ln N	per-period ceiling `eta ln N`
0.30	12	2.4849	0.7455
0.40	6	1.7918	0.7167
0.50	8	2.0794	1.0397

eta	N	ln N	per-period ceiling eta ln N
0.70	4	1.3863	0.9704

Meanwhile **cost does not saturate.** C_t accrues every period whether or not any surplus was delivered. So the honest comparison between an impulse and a cadence is not "same total, same effect." It is "the cadence can deliver T periods of ceiling-limited input while paying T periods of cost; the impulse can deliver *one* period of ceiling-limited input while paying the same T periods of cost." §2.13 computes that comparison for a real case, and the result is not close.

Correction, carried back from Chapter 59. The concession made above is stronger than it looks and this chapter states it plainly rather than leaving it as an aside. Chapter 59 sums Q_t over 400 periods for a 26-unit annual budget delivered three ways at delta = 0.06:

```
impulse      433.3333
weekly       433.3333
four-weekly  433.3333
B/delta   =  433.3333
```

Identical to every digit. Cadence delivers no more integrated stored order than intensity, and any version of the cadence argument resting on the premise that it does is simply false for a linear time-invariant system. Two things do the actual work:

1. **Strict concavity, which needs no ceiling at all.** Theorem 59.1 proves that for a fixed budget T split into n occasions the delivered total D(n) = n g(T/n) is nondecreasing in n whenever g is concave with g(0) = 0 — and Paper A, Theorem 3.3 gives exactly that for f(u) = ln(1 + (N−1)u). Splitting an input across occasions strictly dominates delivering it at once. **This argument is independent of the saturation bound and survives even where the ceiling is never approached.**
2. **The ceiling clips the impulse**, which is the argument this section makes. In Chapter 59's worked annual rite (eta = 0.35, N = 8, per-week ceiling 0.7278), a 26-unit once-yearly burst delivers 0.7278 — **2.8% of the nominal budget. The other 97.2% is not delivered late; it is not delivered at all.**

There is a third difference the infinite integral conceals: the impulse decays to a steady state of exactly 0, while sustained periodic forcing holds m/delta. The integral is equal; the level is not. See Chapter 59, §59.8-§59.11.

2.11 Choosing the period, and the conversions that go with it

The period dt is an operational choice (Paper E, §5.1 lists it among the five required choices). Once made, everything else is expressed in it, and mixing periods is a common and expensive error.

Aggregating up (from short periods to long): retention multiplies.

```
delta_m = 1 − (1 − delta)^m
```

Disaggregating down (from long to short): take the root.

```
delta_1 = 1 − (1 − delta_m)^{1/m}
```

Worked from `delta = 0.05` per week:

horizon `m` (weeks)	combined `delta_m`
1	0.0500
2	0.0975
4	0.1855
13	0.4867
26	0.7365
52	0.9306

And the reverse: an annual `delta` of 0.4596 corresponds to a monthly `delta` of 0.0500 (`1 − (1−0.4596)^{1/12} = 0.05000`).

Design rule R2.2: state the period in the same sentence as the `delta`, always. A bare "`delta = 0.05`" is meaningless. "`delta = 0.05` per week" is a specification.

Design rule R2.3: choose `dt` so that `delta` lands roughly in `[0.02, 0.30]`. Below 0.02 the signal-to-noise on the decay measurement is poor because almost nothing changes per period; above 0.30 the discrete approximation to the underlying continuous process is coarse and the `tau = 1/delta` rule of thumb fails badly (at `delta = 0.5` it overstates the true `tau_exact` by 39%). A system with an eight-month half-life is measured monthly, not daily and not annually.

2.12 Frequency response: why order is low-pass

Evaluate the transfer function on the unit circle, `z = e^{j w}`, and normalise by the DC gain:

```
|H(w)| / |H(0)| = delta / sqrt(1 − 2a cos w + a^2),     a = 1 − delta
```

An input that repeats with period P periods has fundamental frequency `w = 2 pi / P`. The table below gives what fraction of that input's amplitude survives into Q:

input period P	`delta`=0.02	0.05	0.10	0.20	0.50
2	0.0101	0.0256	0.0526	0.1111	0.3333
3	0.0117	0.0296	0.0607	0.1280	0.3780
4	0.0143	0.0362	0.0743	0.1562	0.4472
6	0.0202	0.0512	0.1048	0.2182	0.5774
8	0.0264	0.0669	0.1364	0.2804	0.6786
12	0.0390	0.0986	0.1995	0.3966	0.8069
26	0.0835	0.2081	0.4006	0.6800	0.9465
52	0.1650	0.3910	0.6576	0.8799	0.9857
100	0.3062	0.6325	0.8590	0.9627	0.9961

And the −3 dB cutoff, where the response falls to `1/sqrt(2)` of DC:

`delta`	cutoff `w_c` (rad/period)	period at cutoff
0.01	0.01005	625.2

delta	cutoff w_c (rad/period)	period at cutoff
0.02	0.02020	311.0
0.05	0.05130	122.5
0.10	0.10546	59.6
0.20	0.22408	28.0
0.30	0.36052	17.4
0.50	0.72273	8.7

For small delta, w_c is very nearly delta itself.

Read the tables together and the design consequence is immediate. A weekly system with delta = 0.05 passes 2.6% of a fortnightly oscillation and 63.3% of a two-year one. **Stored order responds to the slow component of what you do and ignores the fast component almost entirely.** The week-to-week variation in how a team feels is invisible to Q. The two-year trend is nearly fully transmitted.

This is why the correct instrument for a low-pass system is a *sustained mean*, not a *peak*. A programme that delivers a variable amount each week, averaging x, produces almost exactly the same Q as one delivering a flat x — because the varying part is filtered out. That is a liberating result for anyone trying to run a real organisation: **you do not have to be consistent week to week; you have to be consistent on the timescale of tau.**

2.13 Ritual as periodic forcing

Take a fixed mean input rate m per period and deliver it in bursts of A = m P every P periods. The steady-state cycle is exact:

```
Q_max = A / (1 − a^P)        immediately after a burst
Q_min = Q_max ∗ a^P          immediately before the next
ripple = Q_max − Q_min = A
```

For delta = 0.10 and m = 0.20 per period:

burst interval P	burst size A	Q_max	Q_min	ripple	ripple as % of mean	mean over cycle
1	0.200	2.0000	1.8000	0.2000	10.0%	2.0000
2	0.400	2.1053	1.7053	0.4000	20.0%	2.0000
3	0.600	2.2140	1.6140	0.6000	30.0%	2.0000
4	0.800	2.3263	1.5263	0.8000	40.0%	2.0000
6	1.200	2.5610	1.3610	1.2000	60.0%	2.0000
12	2.400	3.3446	0.9446	2.4000	120.0%	2.0000

Three things must be said about this table, and the third is the important one.

First, the mean is invariant. Every row has mean 2.0000, equal to m/delta. This is the linear time-invariant result and there is no way around it: *for the same total input, cadence does not change the average stored order.* Anyone who tells you that a weekly ritual produces more *on average* than a quarterly one delivering the same total has not done the arithmetic.

Second, the trough falls with P. At weekly cadence the system never drops below 1.80. At quarterly cadence it drops to 0.94, a 48% swing. If there is any threshold below which the system misbehaves — a level of shared context below which meetings stop working, a level of trust below which people stop volunteering information — the sparse cadence spends time under it while the dense cadence never approaches it. **Averages do not fail systems; minima do.**

Third, and decisively: at realistic parameters you cannot hold m constant while raising P, because A = mP runs into the ceiling eta ln N. This is where the theory bites. The mean is invariant only if the burst can actually be delivered, and Paper D's Corollary 2.2 says it usually cannot.

Take the working example of §2.14: eta = 0.4, N = 6, so the per-week ceiling is eta ln 6 = 0.7167. A four-week burst would need 4 ∗ eta S = 4 ∗ 0.5011 = 2.0044, which is **2.8 times the ceiling**. The burst is clipped to 0.7167. Over a year the weekly cadence delivers 52 ∗ 0.5011 = 26.06 units of input; the four-weekly cadence delivers 13 ∗ 0.7167 = 9.32 units, which is **35.8% as much**. Meanwhile cost is paid all 52 weeks in both cases.

That is the whole argument, and it is arithmetic:

Ritual is a periodic forcing function on a low-pass filter. Cadence beats intensity because surplus saturates at ln N per period while cost does not saturate at all, so the only way to deliver a large total is to deliver it often.

Corrected scope (see §2.10 and Chapter 59). That sentence was originally written as though the ceiling were the whole story. It is half of it. The other half is Theorem 59.1: strict concavity alone makes division across occasions strictly dominant, ceiling or no ceiling. And neither half is "cadence delivers more input" — **it does not; the integrated input is identical.**

Note the corollary, which is unintuitive and useful: **as P shrinks below tau, the ripple vanishes and the ritual becomes indistinguishable from a constant policy.** From §2.12's table, at delta = 0.10 a period-4 forcing transmits only 7.4% of its oscillation. The filter has converted a discrete ceremony into a steady level. **A ritual frequent enough relative to tau stops being an event and becomes infrastructure** — which is exactly what a working ritual feels like from inside, and exactly what Star & Ruhleder (1996) mean when they observe that infrastructure becomes visible only on breakdown.

2.14 The three input regimes, computed

Parameters: delta = 0.08 per week, eta = 0.40, N = 6, u∗ = 0.50, C = 0.15 per week.

```
S = F(u*) = ln(1 + 5 * 0.50) = ln 3.5 = 1.2528
eta S = 0.5011        eta ln N = 0.4 * 1.7918 = 0.7167  (the ceiling)
net weekly drive = eta S - C = 0.3511
Q_inf = 0.3511 / 0.08 = 4.3888
tau_exact = 11.99 weeks,  t_95 = 35.93 weeks,  t_99 = 55.23 weeks
```

Regime A — the impulse. A two-week offsite run at the input ceiling (0.7167 per week), then no further surplus input for the rest of the year. Cost continues at 0.15 per week throughout.

week	0	1	2	3	4	8	12	26	52
Q_t	0.0000	0.5667	1.0881	0.8510	0.6329	0.0000	0.0000	0.0000	0.0000

Peak at week 2, gone by week 8. The offsite is not useless — it puts 1.0881 units of order into a system that had none — but the drive stops and the cost does not, so the system runs down to the admissibility floor within six weeks. **This is the arithmetic of every offsite whose effects "wore off."** They did not wear off. They decayed at `(1−delta)^t` while `C` continued.

Regime B — the sustained step. The same input delivered every week at `eta S = 0.5011`.

week	0	1	2	3	4	8	12	26	52
`Q_t`	0.0000	0.3511	0.6741	0.9713	1.2447	2.1364	2.7752	3.8867	4.3314

Note that Regime B is *behind* Regime A at week 2 (0.674 versus 1.088). At week 4 it has overtaken it. At week 52 it is at 4.3314 against 0.0000. **The regime that looked worse at the first checkpoint is the only one that arrives anywhere.**

Regime C — sparse ritual. An intensive session every fourth week, sized to deliver four weeks of surplus at once — except that it cannot, because 2.0044 exceeds the ceiling 0.7167 and is clipped. Cost continues weekly.

week	0	1	2	3	4	8	12	26	52
`Q_t`	0.0000	0.5667	0.3714	0.1917	0.0263	0.0452	0.0587	0.4393	0.0916

In steady state this regime oscillates between 0.093 and 0.652 with a mean of 0.365 — **8.3% of Regime B's steady state, from a schedule that its participants would describe as "we meet regularly."** The monthly session is genuinely intensive and genuinely valuable; it simply cannot push more than 0.7167 through the ceiling in the week it happens, and the other three weeks pay 0.15 of cost with no input at all.

Summary:

regime	input delivered over 52 weeks	`Q` at week 52	steady-state mean `Q`
A — two-week offsite	1.4334	0.0000	0
B — weekly sustained	26.0575	4.3314	4.3888
C — four-weekly burst	9.3171	0.0916	0.3647

Design rule R2.4: the cadence of a ritual should be short relative to `tau`, and its per-session intensity should be sized to the ceiling `eta ln N`, not to the ambition. Anything beyond the ceiling is discarded in the period it is delivered.

2.15 What a high-pass version would mean

The low-pass character is a consequence of the `Q_t` term appearing on the right-hand side with a positive coefficient. A high-pass version would replace the input with its first difference:

```
Q_{t+1} = (1 − delta) Q_t + eta (S_t − S_{t−1}) − C_t
```

Such a system responds only to *change* in interaction, and is completely blind to a constant level. A steady, high-surplus practice sustained for years would drive it to $-C/\delta$, which is negative: the system would treat a stable excellent state as no input at all.

Systems that behave this way exist, and naming them is useful because they are usually built by accident:

- **News and attention markets.** The stock of attention responds to the derivative of events, not their level. A chronic condition of constant severity generates no coverage.
- **Alarm and exception handling.** By design, an alarm fires on change. This is correct for alarms and catastrophic for anything that must be maintained: a system that reports only exceptions cannot tell you that the ordinary thing is still working.
- **Novelty-driven incentive schemes.** Promotion for launching, none for operating. The organisation's Q receives credit for $S_t - S_{t-1}$ and none for S_t, so maintenance is structurally unrewarded even where everyone agrees it matters.
- **Some reporting regimes.** A dashboard that shows only week-over-week deltas is a high-pass filter installed on top of a low-pass system, and it systematically hides the thing the system actually accumulates.

Social order is low-pass, and it must be, because Q is a stock and a stock integrates. Trust, habit, institutional memory, and furnishing are all quantities that respond to the sustained level of what you do, with the fast fluctuations averaged away. A system whose *measurement* is high-pass while its *dynamics* are low-pass will report noise and miss the trend. This is not a subtle statistical point; it is a mismatch of filter type between the plant and the instrument, and it is one of the most common instrumentation faults in organisational reporting.

Design rule R2.5: match the filter type of the instrument to the filter type of the system. Report levels for stocks and differences for flows. If a dashboard has only difference columns, it cannot measure Q.

2.16 Cascades: when one integrator feeds another

Real systems are often two integrators in series. Individual skill accumulates (δ_1), and team capability accumulates from the flow of trained individuals (δ_2). The cascade is:

$$A_{t+1} = (1 - \delta_1)\, A_t + x_t$$
$$B_{t+1} = (1 - \delta_2)\, B_t + \delta_1 A_t$$

Worked with $\delta_1 = 0.20$, $\delta_2 = 0.08$, unit step into stage 1:

t	stage 1	stage 2
0	0.0000	0.0000
2	1.8000	0.2000
5	3.3616	1.4998
8	4.1611	3.2060
12	4.6564	5.4129
20	4.9424	8.6650
30	4.9938	10.8027
40	4.9993	11.7593
60	5.0000	12.3600

t	stage 1	stage 2
80	5.0000	12.4736

Stage 1 settles at `1/delta_1 = 5.0`; stage 2 settles at `1/delta_2 = 12.5`. Stage 2 reaches half its final value at week 14, 95% at week 43, 99% at week 62 — against 36 and 55 for a single stage with `delta = 0.08` driven directly.

Two facts to carry forward, both checked numerically:

- **The cascade is slower and has an S-shaped rise**, because stage 2 sees nothing until stage 1 has something to give it. The apparent "nothing is happening" phase in months one and two is structural.
- **The cascade does not overshoot.** Two real poles in `(0,1)` give a monotone step response; verified by simulation over 400 periods. Overshoot and ringing require complex poles, which a cascade of leaky integrators cannot produce on its own.

Where does overshoot come from, then? From **feedback with delay** — a controller acting on a stale measurement. Simulating `Q_{t+1} = (1−delta)Q_t + k(target − Q_{t−L})` with `delta = 0.10`, `k = 0.30`, target 5.0:

measurement lag L	peak Q	final Q	overshoot
0	3.7500	3.7500	0.0%
2	4.8826	3.7500	30.2%
4	6.5784	3.7576	75.1%

And at `L = 4` with a more aggressive gain `k = 0.50`, the loop goes unstable outright: the simulated trajectory peaks at 39.2 and then diverges into oscillation with negative excursions, which under the storage law means the admissibility condition has been violated and the model has broken.

Design rule R2.6: the reporting lag of a management control loop must be short relative to `tau`, and the correction gain must fall as the lag rises. An organisation that adjusts quarterly on the basis of annual data is running L of four periods with a large gain, and the resulting whipsaw is not indecision. It is a delayed feedback loop behaving exactly as delayed feedback loops behave.

The Specification

S2.1 — Declare the period. Every parameter shall be stated with its period (`delta = 0.08 per week`). Conversions between periods shall use `delta_m = 1 − (1−delta)^m` and shall be shown in the design record.

S2.2 — Choose the period so that `delta` lands in `[0.02, 0.30]`. Record the justification. If the natural observation cadence forces `delta` outside this band, state the consequence for measurement precision explicitly.

S2.3 — Publish the pole, the time constant, and three settling times. For every modelled subsystem the design record shall contain `1 − delta`, `tau_exact = −1/ln(1−delta)`, `t_50`, `t_95`, and `t_99`, computed from `t_p = ln(1−p)/ln(1−delta)`.

S2.4 — Schedule reviews at or after 3 tau. No initiative shall be evaluated for effect before 3 tau periods have elapsed from the change. Interim checkpoints may verify *inputs* (S, eta, C) but shall not be used to judge Q.

S2.5 — State the DC gain. Record eta/delta, the units of stored order bought per sustained unit of surplus per period, and Q_inf = (eta F(u*) − c*)/delta. Where eta F(u*) < c*, record the negative equilibrium and stop; the regime is unsustainable or the linear model must be replaced (Theorem 3.5).

S2.6 — Size every burst to the ceiling. No scheduled session shall be planned to deliver more than eta ln N of input in one period. Planned excess shall be re-scheduled into additional periods, not compressed.

S2.7 — Specify cadence relative to tau. Ritual interval P shall satisfy P <= tau/2 unless a ripple analysis is recorded showing that Q_min = Q_max (1−delta)^P remains above the system's stated functional threshold.

S2.8 — Record the trough, not only the mean. Any cadence proposal shall report Q_min, Q_max, and ripple, computed from Q_max = A/(1−(1−delta)^P).

S2.9 — Match instrument to plant. Reporting for a stock quantity shall present levels; reporting for a flow quantity shall present rates. A difference-only dashboard is not an acceptable instrument for Q.

S2.10 — Bound the control loop. Where a management loop adjusts inputs on the basis of measured Q, record the measurement lag L in periods and the correction gain k. Require L <= tau/4, and reduce k if L cannot be reduced.

S2.11 — Model-breakdown escalation. If a fitted pole leaves [0,1], or admissibility C_t <= (1−delta)Q_t + eta S_t is violated, the event shall be escalated as a specification failure and re-modelled. It shall not be clipped, smoothed, or averaged away.

The Failure Modes

THE PREMATURE REVIEW — a first-order system is evaluated before 3 tau and judged on the fraction of its step response that has completed.

- *Parameter damaged:* none. The system is fine; the instrument's timing is wrong.
- *Diagnostic signature:* results at review are between 50% and 80% of target and rising; the language used is "slower than hoped" and "losing momentum."
- *Measurement:* compute t_95 = ln(0.05)/ln(1−delta) and compare against the elapsed time. If elapsed < 3 tau, the review is uninformative about Q by construction.
- *Fix:* move the review to 3 tau. Use interim checkpoints on inputs only. Where the sponsor's patience is genuinely shorter than 3 tau, the honest response is to raise delta's counterpart — pick a subsystem with a shorter time constant — not to promise faster arrival.

THE INTENSITY FALLACY — a large one-off input is substituted for a sustained one on the grounds that the total is the same.

- *Parameter damaged:* effectively S (clipped by the ceiling) while C is unaffected.

- *Diagnostic signature:* a peak two to four periods after the event, decay thereafter, and a post-mortem that concludes "we failed to sustain the energy."
- *Measurement:* compare planned session input against `eta ln N`. Compute delivered total across the horizon versus the sustained alternative.
- *Fix:* redistribute. Size sessions to `eta ln N` and increase their frequency until the planned total is deliverable. If the total is not deliverable at any feasible frequency, the total was never available and the plan should say so.

THE SPARSE RITUAL — a ceremony at an interval much longer than `tau`, whose troughs fall below the system's functional threshold even though its mean looks acceptable.

- *Parameter damaged:* none directly; the ripple is the fault.
- *Diagnostic signature:* the ceremony is well-attended and well-regarded, and between occurrences the system behaves as if it does not exist. Participants describe "getting back up to speed."
- *Measurement:* compute `Q_min = Q_max (1−delta)^P` and compare against the threshold below which the system is observed to malfunction. Also compute the delivered total against the ceiling.
- *Fix:* shorten `P` toward `tau/2` or below. Note that halving `P` does not require halving session length: the ceiling caps per-session delivery anyway, so two half-length sessions frequently deliver more than one full-length one.

THE HIGH-PASS INSTRUMENT — the system integrates but the reporting differentiates, so sustained level is invisible and only change is reported.

- *Parameter damaged:* none in the plant; the measurement of `Q` is systematically wrong.
- *Diagnostic signature:* reports consist entirely of week-over-week or quarter-over-quarter changes. A stable, functioning capability generates no report and therefore no resourcing. Maintenance is chronically unfunded (compare Star & Ruhleder 1996: infrastructure becomes visible only on breakdown).
- *Measurement:* audit the reporting pack. Count the fraction of metrics presented as levels versus differences. Ask what report a perfectly stable, perfectly functioning subsystem would generate; if the answer is "none," the instrument is high-pass.
- *Fix:* add level reporting for every stock. Chapter 23 treats the institutional version of this fault at length.

THE OVERDAMPED INSTITUTION — `delta` driven so low that `tau` exceeds the tenure of everyone who could act on a measurement.

- *Parameter damaged:* `delta` (too low).
- *Diagnostic signature:* every proposed change is met with "we tried that" referring to an attempt whose evidence nobody can reconstruct. `t_95` exceeds the average tenure of the decision-makers.
- *Measurement:* compute `t_95` and compare against median tenure in the roles that would act on the result. If `t_95 > tenure`, the institution cannot learn from its own experiments.

- *Fix:* either raise `delta` on the parts of `Q` that should turn over (Chapter 3 §3.12 on deliberate forgetting), or restructure the experiment onto a faster subsystem with a shorter `tau`. Do not attempt to speed up a slow pole by adding input; the pole sets the rate.

THE DELAYED CONTROLLER — corrections applied on the basis of measurements that are several periods stale, producing overshoot and, at higher gain, oscillation.

- *Parameter damaged:* the closed-loop dynamics; the open-loop parameters may be perfectly healthy.
- *Diagnostic signature:* alternating over- and under-correction on roughly a `4L`-period cycle. Hiring surges followed by freezes. Process added then stripped then re-added.
- *Measurement:* record `L` (periods between the state occurring and the decision responding to it) and `k` (fraction of the observed gap corrected per period). Compare `L` against `tau/4`.
- *Fix:* shorten the reporting lag first; reduce the gain second. At `L = 4` and `k = 0.5` the simulated loop diverges outright, so gain reduction alone is not always sufficient.

THE MISMATCHED PERIOD — parameters measured on one period and applied on another without conversion.

- *Parameter damaged:* all of them, silently.
- *Diagnostic signature:* forecasts wrong by a factor that is a power of the period ratio. A monthly-measured `delta` applied as if annual understates decay by a factor near 12 at small `delta` and far more at large `delta`.
- *Measurement:* check every parameter in the record for a stated period. Any parameter without one is suspect.
- *Fix:* restate everything on a single declared period using `delta_m = 1-(1-delta)^m`. Note that `eta` and `C` also carry periods and must be converted; only `u` and `N` are period-free.

Worked Example

A six-category production team choosing its meeting cadence

The system. A production team of nineteen people spanning six genuinely distinct functional categories: design, fabrication, quality, logistics, maintenance, and customer support. Period: one week. The question is whether to replace the current monthly all-hands "deep dive" with a weekly thirty-minute cross-function standup, over the objection that the monthly session is "much more substantial."

Measured parameters.

parameter	value	how obtained
`N`	6	enumeration of functional categories (Chapter 7 protocol)

parameter	value	how obtained
u∗	0.50	disagreement audit and error-decorrelation across category pairs (Chapter 9)
eta	0.40	attributable change in the order proxy per unit surplus, over 14 weeks (Chapter 4)
delta	0.08 per week	log-linear fit on a 5-week zero-input window during a plant shutdown (Chapter 3)
C	0.15 per week	tallied coordination, scheduling, and write-up overhead (Chapter 5)

Derived quantities.

```
ln N        = ln 6                  = 1.7918     hard per-interaction ceiling
N_eff        = 1 + 5 * 0.50            = 3.50       audited effective count
S            = ln 3.50                 = 1.2528     69.9% of ceiling
eta S        = 0.4 * 1.2528            = 0.5011     delivered input per week
eta ln N     = 0.4 * 1.7918            = 0.7167     per-week input CEILING
net drive    = 0.5011 - 0.15           = 0.3511     per week
Q_inf        = 0.3511 / 0.08           = 4.3888     Theorem 3.5
DC gain    = eta/delta = 0.40/0.08   = 5.0      stock per sustained unit of surplus
Thm 3.3 cap = eta ln N / delta         = 8.9588     absolute ceiling on Q
pole         = 0.92
tau_exact    = -1/ln(0.92)             = 11.99 weeks
t_half       = 0.6931 * 11.99          = 8.31 weeks
t_95         = ln(0.05)/ln(0.92)       = 35.93 weeks
t_99         = ln(0.01)/ln(0.92)       = 55.23 weeks
```

The nineteen-person roster contains six categories and buys 3.50 effective ones. The team achieves 69.9% of its theoretical per-interaction ceiling. That is a respectable u∗, and Chapter 9 would say so: the remaining 30% costs more to buy than it returns here, and the binding constraint is elsewhere.

Step 1 — Is the monthly cadence even deliverable?

A monthly session sized to carry four weeks of surplus would need to deliver `4 * eta S = 4 * 0.5011 = 2.0044` in one week. The ceiling is `eta ln N = 0.7167`. **The monthly session is over the ceiling by a factor of 2.80.** No facilitation technique changes this: Paper D, Corollary 2.2 caps `F(u)` at `ln N = 1.7918` and no interaction geometry exceeds it.

So the monthly session actually delivers 0.7167, and 1.288 units of intended input per month — 64.3% of it — do not exist.

Step 2 — Annual delivery, both regimes.

	sessions per year	delivered per session	annual total delivered	cost paid
weekly	52	0.5011	26.0575	52 weeks
monthly (4-weekly)	13	0.7167 (clipped)	9.3171	52 weeks

The weekly cadence delivers **2.80 times** the annual input, and the two regimes pay identical annual cost. The ratio is exactly the clipping factor, which is not a coincidence: it is `4 eta S / eta ln N = 4S/ln N`.

Step 3 — Trajectories.

week	weekly cadence	4-weekly cadence
0	0.0000	0.0000
1	0.3511	0.5667
2	0.6741	0.3714
3	0.9713	0.1917
4	1.2447	0.0263
8	2.1364	0.0452
12	2.7752	0.0587
26	3.8867	0.4393
52	4.3314	0.0916

Steady state: weekly settles at 4.3888. Four-weekly oscillates between 0.0928 and 0.6521 with a mean of 0.3647 — **8.3% of the weekly regime.**

Observe week 1. **The monthly cadence is ahead**, 0.5667 against 0.3511, because a single session at the ceiling puts more in than a single ordinary week. That is the entire empirical basis for the objection that the monthly deep dive is "more substantial," and it is correct — for one week. By week 2 the weekly regime is ahead and it never trails again.

Step 4 — Ripple analysis for the sparse regime, per S2.8.

For `P = 4`, `a = 0.92`, `A = 0.7167`:

```
a^P    = 0.92^4                        = 0.7164
Q_max  = A/(1 - a^P) = 0.7167/0.2836   = 2.5271   (if cost were zero)
Q_min  = Q_max * a^P                   = 1.8104   (if cost were zero)
```

With cost at 0.15 per week the simulated cycle is far lower — 0.0928 to 0.6521 — because the weekly cost draw consumes most of what the burst delivers before the next burst arrives. **The ripple formula is the no-cost idealisation; the real trough is set by the cost drawn between bursts, and that is the number that must be reported.**

Step 5 — Filter check, per §2.12. At `delta = 0.08`, an input oscillating with period 4 weeks is transmitted at roughly 6% of its amplitude (interpolating between the `delta = 0.05` value of 0.0362 and the `delta = 0.10` value of 0.0743). The team's stored order is essentially blind to the monthly rhythm as *rhythm*; what reaches `Q` is the mean, and the mean is what clipping has destroyed.

Step 6 — Review scheduling, per S2.4. `3 tau = 35.98` weeks. The correct evaluation point for the weekly-cadence change is **week 36**, when the system will be at 95.03% of `Q_inf`. A review at week 12 will find 63.2%, and a review at week 26 will find 88.6%. Both are on the curve. Neither is evidence of a problem. The design record states all three numbers in advance so that nobody has to argue about it later.

Step 7 — Headroom check, per Chapter 6. At the new steady state `Q_inf = 4.3888`:

`h = eta ln N − delta Q − C = 0.7167 − 0.08 ∗ 4.3888 − 0.15 = 0.2155`

Positive, so growth remains available, and `r_max = 0.2155` per week is the largest increment that may be promised (Corollary 4.2). Headroom reaches zero at `Q = (0.7167 − 0.15)/0.08 = 7.0838`, safely above the operating point.

At the *old* monthly cadence with Q near 0.36, headroom was `0.7167 − 0.08(0.36) − 0.15 = 0.5379` — apparently *higher*. This is the trap Chapter 6 names: **headroom is high in a system that has stored nothing, and that is not a sign of health.** A team with no accumulated order has enormous room to build and nothing built. Headroom is a budget, not a score.

Step 8 — Recommendation and falsification.

Adopt the weekly cadence. Retain a quarterly longer session for the genuinely serial work that cannot be decomposed into weekly increments, sized to the ceiling and understood to contribute `0.7167` in its week rather than a quarter's worth.

Predicted: Q rises from a mean of 0.365 to 4.389, reaching 95% of the new value at week 36.

Preregistered falsification, per Chapter 25: fit `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t` on weeks 1-36 with `0 <= a <= 1` and `b >= 0`, freeze the coefficients, and report on weeks 37-72 against four baselines — persistence, direct-input, matched autoregressive, and shuffled-input control. The prediction is falsified if:

1. the fitted a falls outside `[0,1]` (the first-order form is wrong);
2. the fitted b does not exceed its pre-intervention value (the cadence change did not raise delivered input);
3. the frozen model fails to beat the **shuffled-input control** on held-out data — which would mean that the *timing* of interaction carries no information beyond its volume, and the entire cadence argument in this chapter would not apply to this system.

The shuffled-input control is the decisive test for this chapter's thesis. Every claim made here about cadence is a claim that *when* surplus arrives matters, not merely how much. If shuffling the input sequence does not degrade the fit, the low-pass structure is not present in this system and the chapter's design rules do not apply to it. Chapter 25 explains why this control is the one most often omitted, and this is precisely why omitting it is not acceptable.

What this chapter has established

The storage law is a first-order discrete low-pass filter with pole `1 − delta`, time constant `1/delta` periods, DC gain `eta/delta`, and steady state `Q_inf = (eta F(u∗) − c∗)/delta`, which is Paper E's Theorem 3.5 read as a DC gain.

Its step response is monotone and reaches 95% in about `3 tau` and 99% in about `5 tau`, which sets when an initiative may honestly be reviewed. Its impulse response is pure decay, which is why one-off events fade at a rate you can compute in advance. Its frequency response attenuates everything fast compared to `tau`, which is why week-to-week variation is invisible to stored order and why a sufficiently frequent ritual stops being an event and becomes infrastructure.

And because surplus saturates at `ln N` per period while cost does not saturate at all, the total a system can deliver is set by how often it delivers, not by how hard it delivers on any one occasion. Two qualifications, both established in §2.10 and both carried back from Chapter 59: the *integrated* effect of an impulse and of the same total delivered periodically is identical to every digit, so cadence never wins by delivering more; and the ceiling is only half the reason it wins anyway, the other half being the strict concavity of `f` (Theorem 59.1), which makes division across occasions dominant even where nothing is ever clipped.

Ritual is a periodic forcing function on a low-pass filter. That is why cadence beats intensity, and the margin can be computed before anyone commits to a calendar.

Chapter 3 takes the pole itself as the object of engineering: what `delta` is made of in real systems, how to measure it, how to drive it down where retention is the goal — and where you should deliberately drive it up, because a system that cannot forget cannot change its mind.

Chapter 3 . Loss — Engineering Delta

`delta` is the parameter that makes the whole machine work, and it is the one that no organisation has ever put on a dashboard.

Chapter 1 showed why it is load-bearing: `delta > 0` makes the update map a contraction, and contraction is what supplies a destination, a destination independent of history, and error decay. Chapter 2 showed what it controls: the pole, the time constant, the settling time, the DC gain, and the cutoff frequency. Every one of those is `delta` wearing a different hat.

This chapter takes `delta` as an object of engineering. Four questions, in order:

1. **What is it made of?** Loss is not a single phenomenon; it is a set of named channels that compose in a specific way.
2. **How do you measure it?** With a zero-input window and a log-linear fit, and with a specific list of things that will contaminate the measurement.
3. **How do you drive it down?** The artifact ladder, and the two prices you pay for descending it.
4. **When should you drive it up?** More often than anyone expects, because **a system with no forgetting cannot learn.**

The last question is the one this chapter exists for. Most practitioners, once they understand `delta`, immediately try to minimise it everywhere. That is a mistake with a precise cost, and §3.12 computes it.

The Law

3.1 Where `delta` sits

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     delta in [0,1]
```

`delta` is the fraction of stored order lost per period **in the absence of maintenance**. That qualifier is the whole definition. `delta` is not "how much we lose"; it is "how much we would lose if we did nothing," and the maintenance we actually perform appears elsewhere in the equation — partly in `eta S_t` where maintenance regenerates structure, and partly in `C_t` where the effort of maintaining is paid for.

Confusing `delta` with observed net decline is the most common estimation error in this book, and §3.7 gives the procedure that avoids it.

3.2 What the theorems say about `delta`

Everything, essentially. Reading Paper E §3 as a set of statements about `delta`:

Theorem 3.1 (exact solution): `Q_t = (1−delta)^t Q_0 + sum_{k=0}^{t−1} (1−delta)^{t−1−k}(eta S_k − C_k)`. `delta` sets the discount rate on history. The weight on an input `j` periods old is `(1−delta)^j`.

Corollary 3.2 (one-step growth): `Q_{t+1} > Q_t` iff `eta S_t > delta Q_t + C_t`. The term `delta Q_t` is the **maintenance load**, and it is proportional to how much you have.

Theorem 3.3 (uniform bound): `Q_t <= max{Q_0, eta ln N / delta}`. `delta` is in the denominator of the ceiling. Halving `delta` doubles the maximum attainable stock. **If `delta = 0`, the bound degrades to `Q_t <= Q_0 + t eta ln N` — linear in `t` and unbounded.**

Corollary 3.4 (sustainability ceiling): a nondecreasing step requires `C_t <= eta ln N − delta Q_t`. As `Q` grows, `delta Q` eats the cost budget, and past `Q = (eta ln N − C_t)/delta` **no choice of interaction geometry can prevent order from decreasing.**

Theorem 3.5 (steady state): with `delta > 0`, `Q_inf = (eta F(u∗) − c∗)/delta`. `delta` is the denominator of the fixed point. If `eta F(u∗) < c∗` the equilibrium is negative and the theorem declares its own failure: the regime is unsustainable or the linear model must be replaced.

Headroom: `h_t = eta ln N − delta Q_t − C_t`. `delta` appears here as the coefficient on the maintenance term, which is Chapter 6's entire subject.

Banach: the contraction modulus is exactly `k = 1 − delta`. `delta > 0` is necessary and sufficient for the map to be a contraction on the real line, and therefore for existence, uniqueness, and global convergence.

3.3 Half-life

The transient `(1−delta)^t` is a geometric decay, so it has a half-life:

```
t_half = ln 2 / (−ln(1 − delta))
```

This is the single most useful derived quantity in the manual for talking to non-specialists. "Our `delta` is 0.10 per week" means nothing to a director. "Half of everything we know how to do, that isn't written down or exercised, is gone in about six and a half weeks" is a sentence that gets a budget.

Computed for the range that matters:

`delta` (per period)	half-life (periods)	to 10% remaining	to 1% remaining	retained after 10 periods
0.01	68.97	229.11	458.21	0.9044
0.02	34.31	113.97	227.95	0.8171
0.05	13.51	44.89	89.78	0.5987
0.10	6.58	21.85	43.71	0.3487
0.20	3.11	10.32	20.64	0.1074
0.30	1.94	6.46	12.91	0.0282
0.50	1.00	3.32	6.64	0.0010

Note the nonlinearity. Going from `delta = 0.10` to `delta = 0.05` roughly doubles the half-life (6.58 to 13.51) and doubles `Q_inf`. Going from 0.05 to 0.02 multiplies it by 2.54. **Each halving of `delta` doubles both the ceiling and the settling time**, and the second of those is the price nobody budgets for.

3.4 Honesty clauses

- `delta` is **period-dependent**. It is meaningless without a stated period (Chapter 2, S2.1).

- `delta` measured on a proxy Y is the decay rate of the proxy. **You must justify that Y is not a repackaging of the input U** (Paper E, §6.3), or you are measuring the decay of your own measurement instrument.
- The exponential decay form is a **modelling assumption, badged**. Real loss processes can be non-exponential — a cliff at a single departure, a step at a reorganisation, a long tail from written records. Fitting an exponential to a cliff produces a number that is wrong in a specific and predictable direction, and §3.7 says how to detect it.
- Every `Q_inf` computed from an estimated `delta` inherits its uncertainty **in the denominator**, which means small errors in `delta` produce large errors in `Q_inf`. A 20% error in `delta` produces a 25% error in `Q_inf`. Report ranges, not points.
- As deployed the surplus functional is a **lower bound** on the ensemble's information (Paper A, Paper D). Everything downstream is a floor.
- **No thermodynamic claim.** `delta` is not entropy production and Paper E proves no negentropy law. It is a coefficient in a difference equation.

The Engineering

3.5 What `delta` is made of

Loss is not one thing. In any real system it decomposes into named channels, each with its own mechanism, its own measurement, and its own available intervention. The eight that recur:

(1) Procedural forgetting. An individual's memory of a procedure they have not executed recently. Fast for multi-step procedures with arbitrary ordering; slow for procedures with strong internal logic. This is the channel that quarterly exercises address.

(2) Turnover. A departure removes whatever fraction of `Q` lived in that person and nowhere else. §3.10 treats this as its own subject because it is usually the largest single channel and the easiest to compute.

(3) Documentation entropy. Written records do not decay by fading; they decay by **becoming false**. The document is intact and the world has moved. Documentation entropy is the reason a written artifact does not reduce `delta` to zero, and it is why *executable* documentation — a test, a script, a check that fails when it is wrong — has a materially lower `delta` than prose.

(4) Environment drift. The system the stored order refers to has changed. A runbook that references a decommissioned server, a floor plan that predates the renovation, a policy that names a department that no longer exists. This channel is invisible until exercised, which is what makes it dangerous.

(5) Relationship decay. Trust and working familiarity between two people, absent interaction. This is a genuine `delta` channel and it is one of the few whose maintenance cost is almost pure `C` — a standing one-to-one produces very little `S`, and its purpose is not to produce `S` but to suppress `delta`.

(6) Skill atrophy. Distinct from procedural forgetting: the procedure is remembered and the execution has degraded. Motor skill, judgement calibration, and fluency all sit here.

(7) Physical wear. Furniture, buildings, instruments, the acoustic treatment that has come loose. The most literal `delta` and the only one with an established maintenance profession attached,

which is instructive: the reason building maintenance is a discipline and institutional maintenance is not is that physical `delta` is visible.

(8) Context loss. The rule survives and its rationale does not. **This channel is worse than plain forgetting**, because it leaves behind a rule that is still enforced and can no longer be evaluated. A rule whose rationale has decayed cannot be retired, because nobody can establish that it is safe to retire it, so context loss converts ordinary `Q` into permanent `C`.

Channel (8) deserves the emphasis. Ordinary loss removes structure and frees the budget it was consuming. Context loss removes the *justification* while leaving the structure, and the structure keeps drawing cost forever. **Context loss is the mechanism by which decay produces ossification rather than clearance**, and it is the reason Chapter 23's discussion of deliberate shedding has to be a design activity rather than a natural process.

3.6 Channels compose multiplicatively

Retention multiplies; loss does not add. If k independent channels each retain a fraction `1 − delta_i` per period, then

`1 − delta_total = product_i (1 − delta_i)` so `delta_total = 1 − product_i (1 − delta_i)`

Worked, for four channels in a team:

channel	`delta_i`
turnover	0.030
procedural forgetting	0.050
documentation rot	0.020
tool churn	0.015
naive sum	**0.1150**
true combined `1 − prod(1−delta_i)`	**0.1105**

The naive sum overstates by 0.0045, about 4% relative. At these magnitudes the difference is small and the sum is a serviceable approximation. It stops being serviceable quickly: four channels at 0.20 each give a naive sum of 0.80 and a true combined `delta` of 0.5904, an overstatement of 35%.

Design rule R3.1: use the product form whenever any channel exceeds 0.10, and always in the design record. The sum is acceptable only as a back-of-envelope check at small `delta`.

The corollary matters more than the arithmetic. **The channels are not fungible, and reducing the largest one has the largest effect on the product.** Halving the 0.050 forgetting channel to 0.025 takes the combined `delta` from 0.1105 to 0.0871, a 21.2% reduction. Halving the 0.015 tool-churn channel takes it only to 0.1037, a 6.2% reduction. Attack the biggest channel; the multiplicative structure rewards it.

3.7 Measuring `delta`: the zero-input window

The procedure.

1. **Choose the order proxy Y.** State what it counts. Write the justification that Y is not a repackaging of the input U (Paper E, §6.3). Have someone else check it.
2. **Find or create a window in which S contributions to Y are near zero.** Candidates: a plant shutdown; a hiring or feature freeze; a holiday period; the interval after a contributor leaves and before a replacement starts; a decommissioned tool still nominally in service; a documented practice that nobody has exercised.
3. **Measure Y at several points in the window.** Two points give a number; five give a fit and a residual, and the residual is where the information about model form lives.
4. **Fit log-linearly.** Regress `ln Y_t` on `t`. The slope is `ln(1−delta)`, so `delta = 1 − exp(slope)`. Report `R^2`.
5. **Convert to your working period** if the window was measured on a different one: `delta_1 = 1 − (1−delta_m)^{1/m}`.
6. **Sanity-check against the half-life.** Does `ln 2 / (−ln(1−delta))` match what practitioners say when asked "how long before this would be gone if we stopped?" A factor-of-two disagreement is normal; a factor-of-ten disagreement means the proxy is wrong.

The two-point estimator, for when that is all you have:

```
delta = 1 − (Y_k / Y_0)^{1/k}
```

The pitfalls, each with its detection:

- **Contaminated window.** Input leaked in. Detection: the fit shows a plateau or an upturn, or `R^2` is poor with structured residuals. Fix: shorten the window or find a cleaner one.
- **The proxy is the input.** Y is really a measure of U or S, so its "decay" is just the activity stopping. Detection: Y drops to near zero the moment the window opens, rather than decaying. This is Paper E §6.3's warning and it is the most common single error in the field.
- **Cliff, not exponential.** The loss was one departure, not a rate. Detection: residuals show a single large step; the per-step estimates are wildly heterogeneous. Fix: model the departure as an event and estimate `delta` from the remainder.
- **Censoring.** The proxy has a floor — you cannot count fewer than zero runbooks, and the last few are the most robust. Detection: the fit flattens at low Y. Fix: fit only the upper range and say so.
- **Reactive maintenance.** People notice Y falling and intervene, which is `eta S` sneaking back in. Detection: ask. Fix: instrument the window so intervention is recorded.

Worked estimation. A monthly proxy over a six-month zero-input window:

month `t`	0	1	2	3	4	5	6
`Y_t`	100	92	84	79	71	66	60

Per-step naive estimates `1 − Y_t/Y_{t−1}`: 0.0800, 0.0870, 0.0595, 0.1013, 0.0704, 0.0909. The spread is 0.0595 to 0.1013, which is what single-step estimates always look like and why nobody should use one.

Endpoint two-point estimator: `1 − (60/100)^{1/6} = 0.0816`.

Log-linear regression: slope `−0.08446`, so **`delta = 0.0810`**, fitted `Y_0 = 100.16`, **`R^2 = 0.9980`**. Half-life `ln 2 / (−ln 0.9190) = 8.21 months`.

The regression and the two-point estimator agree to within 0.8% here, which is a good sign — the decay really is close to exponential over this range. **When they disagree by more than about 15%, believe neither and look at the residuals**, because the disagreement means the shape is not geometric and the whole exponential model is in question for this proxy.

3.8 What if you cannot get a zero-input window?

Sometimes you cannot stop the system and cannot find a natural gap. Then `delta` must be estimated jointly with `eta` from the full model, which is Chapter 25's procedure:

```
Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t,    0 <= a <= 1,  b >= 0,  a = 1−delta,  b = eta
```

Fit on training data, freeze the coefficients, report held-out performance against the four required baselines — persistence, direct-input, matched autoregressive, and shuffled-input control.

Two cautions. First, `a` and `b` are correlated in the fit: "high retention, low conversion" and "low retention, high conversion" describe the data almost equally well, and separating them requires variation in `U` that the data may not contain. Second, **a fitted `a` outside `[0,1]` is not a number to be clipped; it is a specification failure** and must be escalated as such (Chapter 2, S2.11).

Prefer the zero-input window when you can get one. A single documented shutdown identifies `delta` directly and is worth more than a year of regression on confounded data.

3.9 Periods, and the conversions

`delta` is a per-period quantity, so it must always be reported with its period. Aggregation and disaggregation:

```
delta_m = 1 − (1−delta_1)^m          (short period to long)
delta_1 = 1 − (1−delta_m)^{1/m}      (long period to short)
```

Monthly to annual, computed:

monthly `delta`	annual retention `(1−delta)^12`	annual `delta`
0.01	0.8864	0.1136
0.02	0.7847	0.2153
0.05	0.5404	0.4596
0.10	0.2824	0.7176
0.15	0.1422	0.8578
0.20	0.0687	0.9313
0.30	0.0138	0.9862

And back the other way:

annual `delta`	implied monthly `delta`
0.05	0.0043
0.10	0.0087
0.20	0.0184

annual `delta`	implied monthly `delta`
0.30	0.0293
0.50	0.0561
0.70	0.0955

A monthly `delta` of 0.05 — which sounds modest, and which a manager would describe as "we lose a bit of ground each month" — destroys 46% of unmaintained stored order in a year. That single row is the most persuasive line in most `delta` presentations, and it is arithmetic, not rhetoric.

3.10 Turnover as a `delta` term

Turnover is usually the largest channel and it has the cleanest formula. If a fraction p of members leave per period, and a fraction `phi` of stored order is **person-borne** — held in individuals and nowhere else — then to first order

```
delta_turnover = p * phi
```

The two factors are independently controllable, and this is the practical point. You can reduce turnover (expensive, partly outside your control, and sometimes undesirable — see §3.12) or you can reduce the person-borne fraction (cheap, entirely within your control, and beneficial for other reasons). **`phi` is the parameter, not `p`.**

Computed, at 20% annual turnover (monthly `p = 0.0184`):

person-borne fraction `phi`	monthly `delta_turnover`	annual equivalent
0.90	0.0166	0.1818
0.75	0.0138	0.1538
0.60	0.0111	0.1249
0.40	0.0074	0.0849
0.20	0.0037	0.0433
0.05	0.0009	0.0110

Moving `phi` from 0.90 to 0.20 cuts the turnover channel by a factor of 4.5 without changing a single thing about retention or recruitment. Every mechanism in Chapter 4 for raising `eta` also lowers `phi`, because they are the same mechanisms: writing it down, building the object, training the second person, changing the default. This is the strongest instance of a general pattern — **the artifact that captures surplus is the same artifact that survives the person who made it.**

Conway's Law (1967) enters here with a specific quantitative consequence. Systems mirror the communication structure of the organisation that builds them, which means a substantial part of Q is encoded in the org chart rather than in any document: who talks to whom, which pairs have a working channel, where the informal escalation paths run. **A reorganisation is therefore a `delta` event, and a large one**, because it invalidates the part of stored order that lived in the communication topology. The reorganised system retains its documents and loses its routing.

This is measurable. Before a reorganisation, sample the "who would you ask about X" network for a set of X. After, sample it again. The fraction of queries whose answer changed, or became

"I don't know," is a direct estimate of the reorganisation's one-off `delta` contribution. It is almost never measured, and it is almost always large.

Brooks (1975) supplies the other half: adding people to a late project makes it later, because the new arrivals must be brought up on stored order that is largely person-borne, and the bring-up consumes the time of exactly the people who hold it. In the law's terms, onboarding is a `C` term proportional to `phi Q`, and it is charged against the same budget as everything else.

3.11 Engineering `delta` down

The intervention is always the same in form: **move stored order down the artifact ladder**, from media with high `delta` toward media with low `delta`.

```
 THE ARTIFACT LADDER            typical delta        why
 -------------------            (per month, rough)
one person's memory           0.10 – 0.30        forgetting + they may leave
two people's shared memory    0.06 – 0.20       redundancy halves the departure risk
spoken transmission / story   0.05 – 0.15       survives departure, drifts in content
written prose document       0.02 – 0.08       survives; goes silently false
 diagram / checklist             0.02 – 0.06          harder to be ambiguous
executable check / test       0.005 – 0.02       FAILS when it becomes false
physical object / built form  0.001 – 0.02       must be actively removed to be lost
trained cohort (>3 people)    0.005 – 0.03       regenerates itself if cadence holds

 These bands are ORDER-OF-MAGNITUDE ORIENTATION, badged as such. They are
 not measurements. The point of this chapter is that you measure your own.
```

That badge is not decorative. The bands above are the author's working priors for structuring a survey, not findings. **Any number in your design record must come from your own zero-input window, not from this ladder.** The ladder tells you which direction to move and roughly how much to expect; §3.7 tells you what is actually true in your system.

The rungs, with what each actually buys:

Write it down. The single largest step on the ladder, and the cheapest. It converts a channel that includes departure risk into one that does not. It does *not* eliminate documentation entropy — the document goes false — so it lowers `delta` by roughly a factor of 3 to 5, not to zero.

Make it executable. The step from prose to a check that fails is the largest reduction available short of building a physical object, and it is available in far more domains than people assume: automated tests, monthly restore drills, a checklist with a signature line, a form that will not submit if a field is inconsistent. **Executable documentation is the only kind that reports its own decay.** This is the direct answer to documentation entropy: prose rots silently, and a test that rots turns red.

Build the object. A changed floor plan, a purpose-built jig, a room with the acoustic treatment installed. Physical structure has the lowest `delta` available because losing it requires deliberate removal. Parts III of this manual is largely an application of this rung.

Train the second, third, and fourth person. Redundancy in people is `phi` reduction, and past about three holders it becomes self-regenerating: the cohort can rebuild a departed member's share without external input. **Below three holders, a practice is a person with a job title.**

Institute the ritual. Chapter 2 established that a periodic forcing function at interval `P <= tau/2` becomes indistinguishable from a constant policy. In `delta` terms, exercise suppresses procedural forgetting and environment drift simultaneously: the exercise *reveals* the drift, which is the only way drift is ever revealed.

Apprenticeship. The transfer of context, not just procedure. This is the only mechanism this manual knows of that addresses channel (8), context loss, because rationale is not writable in a form anyone reads. Apprenticeship is expensive and slow and has no substitute.

Design for patina. Jacobs (1961) observed that aged buildings carry a mix of uses new construction cannot, partly because a building that ages *well* accumulates adaptation rather than requiring replacement. **A material that looks better worn has a lower effective `delta` than one that looks worse worn**, because the replacement decision is deferred. Chapter 16 develops it.

3.11.1 The two prices of a low `delta`

Both are routinely omitted from proposals, and both are computable.

Price one: `C` rises. Every rung of the ladder costs something to maintain. Tests must be kept green; documents must be reviewed; drills must be run; the second and third holders must spend time holding. Since `C` enters the sustainability condition linearly (Corollary 3.4: `C_t <= eta ln N − delta Q_t`) and `delta` enters multiplied by `Q_t`, the trade is favourable only while `delta Q_t` saved exceeds `C` added. **At small `Q`, lowering `delta` is not worth its cost.** A three-month project should not build institutional memory; it should finish.

Price two: the system gets slower. This one is invisible and it is the more dangerous of the two. From Chapter 2, `t_95 = ln(0.05)/ln(1−delta)`:

`delta`	half-life	`t_95`	`Q_inf` at drive 7.775
0.0810	8.21	35.5	96.0
0.0405	16.76	72.5	191.9
0.0368	18.49	79.9	211.3

Halving `delta` doubles `Q_inf` and doubles `t_95`. **You have built a system that holds twice as much and takes twice as long to change its mind about anything.** If the environment your system responds to moves faster than the new `t_95`, you have bought stability you cannot use and lost adaptability you needed. Chapter 2's THE OVERDAMPED INSTITUTION is this failure at its limit: when `t_95` exceeds the tenure of the decision-makers, the institution cannot learn from its own experiments.

Design rule R3.2: state the new `t_95` in the same paragraph as the new `Q_inf`, every time a `delta` reduction is proposed.

3.11.2 The invisibility problem

Star & Ruhleder (1996): infrastructure becomes visible only on breakdown. This has a precise consequence in this model, and it is structural rather than cultural.

Maintenance is the activity that suppresses `delta`. Its output is a *non-event*: the thing that did not go wrong. A high-pass reporting instrument (Chapter 2, §2.15) generates no signal from a

non-event, so maintenance produces no visible result while it is working and produces a highly visible result the moment it stops. **The budget-allocation process is therefore systematically biased against exactly the activity that keeps `delta` low**, and the bias does not require anyone to be foolish; it follows from the filter type of the instrument.

Three counters, all instrumentation rather than exhortation:

1. **Report `delta` itself as a headline metric.** Not incidents; the decay rate. It is a level, not a difference, and it is visible when things are fine.
2. **Report the counterfactual maintenance load `delta Q_t`** — the per-period cost that would be incurred with no maintenance. It makes the invisible work quantitative.
3. **Run scheduled failure drills.** A drill converts a non-event into an event on a schedule you control — executable documentation applied to organisational rather than technical structure.

3.12 When you want `delta` high

Here is the half of the chapter that most `delta` discussions omit entirely.

`delta` is a per-component parameter, not a global constant. **Different parts of `Q` should have radically different `delta`, and setting them all low is a design failure, not a design goal.**

What should decay fast:

- **Grudges and interpersonal records.** A stored account of who wronged whom is `Q` in the model's terms — it is durable structure that shapes future interaction — and its persistence drives u toward zero within the affected pair by making the interaction predictable and defensive. High `delta` here is called forgiveness, and it is an engineering choice with a half-life.
- **Obsolete procedure.** A rule whose original condition no longer obtains draws `C` forever and produces nothing. See channel (8): once its rationale has decayed, it cannot be retired by argument, only by scheduled expiry.
- **Bad habit.** Any practice that made sense once. The whole point of a practice is that it does not require re-deciding, which is exactly what makes an obsolete one durable.
- **Roadmaps, plans, and forecasts.** These are supposed to decay. A two-year-old roadmap that is still being executed is a `delta`-too-low failure, and it is common.
- **Metrics that have become targets.** Goodhart's Law and Campbell's Law both say a measure that becomes a target ceases to be a measure. In this model, a metric is stored order with a `delta`, and **the correct response to Goodhart capture is to raise the metric's `delta` — retire it and replace it — rather than to defend it.** A permanent metric is a guaranteed-corrupted metric.
- **Reputational classification of people.** Any stored judgement about a person's capability becomes self-confirming through assignment. High `delta` on these records is what makes development possible.

The general statement, and it is the load-bearing sentence of this chapter: a system with no forgetting cannot learn.

The proof is one row of a table. An error injected at `t = 0` retains weight `(1−delta)^t`:

periods since the error	delta=0	0.01	0.05	0.10	0.20
1	1.00000	0.99000	0.95000	0.90000	0.80000
5	1.00000	0.95099	0.77378	0.59049	0.32768
10	1.00000	0.90438	0.59874	0.34868	0.10737
25	1.00000	0.77782	0.27739	0.07179	0.00378
50	1.00000	0.60501	0.07694	0.00515	0.00001
100	1.00000	0.36603	0.00592	0.00003	0.00000

Learning is the replacement of an old weighting with a new one. Replacement requires that the old weighting fall. **At `delta = 0` no weighting ever falls, so no replacement ever completes**, and what looks like learning is only accumulation: the new belief is added alongside the old one, both at full weight, and the system now holds a contradiction it cannot resolve.

This is not a metaphor about organisational culture. It is the first column of the table.

Note also what high `delta` does at the *other* end of the ladder. A component with `delta = 0.25` has a half-life of 2.41 periods and 0.1% residue after 24 periods. That is a component which can be replaced wholesale without any removal effort at all: **you retire it by declining to maintain it.** Deliberate high-`delta` design — sunset clauses, term limits, expiring authorisations, dated documents that go stale visibly — is the cheapest removal mechanism that exists, because it costs nothing to invoke. Chapter 23 develops the institutional form.

3.13 The `delta` portfolio

Design output for any real system is not a `delta`; it is a table of `delta`s, one per component, each justified.

component	target `delta` (per month)	half-life (months)	rationale
constitutional / safety-critical rules	0.002	346.2	must survive every individual; changing them must be deliberate and slow
safety procedures	0.010	69.0	durable, but must be re-derivable and periodically re-examined
core tooling and standards	0.050	13.5	durable within a technology generation; must not outlive it
team practice and convention	0.100	6.6	should be re-decided roughly twice a year by exercise
roadmap and plan	0.250	2.4	must go stale visibly; a stale plan being executed is the failure

component	target delta (per month)	half-life (months)	rationale
interpersonal records / grudges	0.400	1.4	near-total clearance in six months by design

Two design rules come straight out of this table.

Design rule R3.3: no component gets `delta = 0`. Even the constitutional layer is 0.002, which is a 346-month half-life and an amendment process, not a stone tablet. A component at exactly zero has no fixed point, no error decay, and no mechanism of correction (Chapter 1, §1.8).

Design rule R3.4: `delta` should decrease monotonically with the component's intended lifetime, and every component must have a stated intended lifetime. A component nobody has assigned a lifetime to will default to `delta = 0` through neglect, because retirement requires a decision and persistence requires none.

3.14 The trade computed: `delta` down versus `eta` up

Which lever is cheaper? Chapter 26 ranks them formally; here is the arithmetic that makes the ranking concrete.

Take `eta = 0.5`, `N = 8`, `u* = 0.6` so `S = 1.6487`, `C = 0.20`, and `delta = 0.20`:

```
baseline:             Q_inf = (0.5 * 1.6487 - 0.20)/0.20 = 3.1216
halve delta to 0.10:  Q_inf = (0.5 * 1.6487 - 0.20)/0.10 = 6.2433    (+100.0%)
double eta to 1.0:    Q_inf = (1.0 * 1.6487 - 0.20)/0.20 = 7.2433    (+132.0%)
```

Doubling `eta` beats halving `delta`, and it beats it by 32 percentage points. The reason is structural: `eta` multiplies S before the cost is subtracted, so doubling it adds `eta S = 0.8243` to the numerator, whereas halving `delta` merely doubles whatever the numerator already is — including the drag of `–C`. The larger `C` is relative to `eta S`, the more this asymmetry favours `eta`.

That is the general case and Chapter 4 will lean on it. But note the exception, because it is important and it is where this chapter earns its place in Part I: **`delta` reduction is the better lever whenever the numerator `eta F(u*) – c*` is already large and the problem is that nothing stays.** A system generating excellent surplus, converting it well, and losing it all by the next quarter has a `delta` problem, and no amount of additional `eta` will fix a leak. The diagnostic is simple: **if `Q` is oscillating or flat while `eta S` is comfortably above `C`, the fault is `delta`.** §3.7's zero-input window will confirm it in one measurement.

The Specification

S3.1 — `delta` shall be measured, not assumed. Every design record shall contain a `delta` estimate obtained from a documented zero-input window or from a frozen, held-out fit per Chapter 25. Orientation bands (§3.11) are not measurements and shall not be cited as such.

S3.2 — State the period. `delta` is reported as `delta = X per <period>`, always. Conversions use `delta_m = 1–(1–delta_1)^m` and are shown.

S3.3 — Decompose into named channels. The estimate shall be accompanied by an enumeration of loss channels with an individual estimate for each, combined by `delta_total = 1 − prod(1−delta_i)`. The residual between the combined channel estimate and the measured total shall be reported. A residual above 25% of the measured total means the enumeration is incomplete.

S3.4 — Report the half-life. Every `delta` shall be reported alongside `t_half = ln 2 / (−ln(1−delta))` in the same line.

S3.5 — Justify the proxy. Written justification that the order proxy `Y` is not a repackaging of the input `U` (Paper E, §6.3) shall precede any `delta` estimation.

S3.6 — Report `phi`. The person-borne fraction of stored order shall be estimated and reported. Turnover contribution is recorded as `delta_turnover = p * phi` with both factors stated separately.

S3.7 — Treat reorganisation as a `delta` event. Any structural reorganisation shall be preceded and followed by a "who would you ask about X" network sample over a fixed question set, and the fraction of changed or lost answers recorded as a one-off `delta` contribution.

S3.8 — Every component has an intended lifetime and a target `delta`. The design record contains a `delta` portfolio table (§3.13). **No component is assigned `delta = 0`.**

S3.9 — State both prices of a `delta` reduction. Any proposal to lower `delta` shall state, in the same section: the added `C` per period, and the new `t_95 = ln(0.05)/ln(1−delta_new)`.

S3.10 — Executable over prose where available. Where stored order can be expressed as a check that fails when it becomes false, it shall be, and the prose version shall be derived from the check rather than maintained in parallel.

S3.11 — Minimum three holders. No practice on which the system depends shall have fewer than three trained holders. Where fewer, the shortfall is recorded as an explicit `phi` risk with its `delta_turnover` contribution computed.

S3.12 — Maintenance shall be reported as a level. `delta` and the counterfactual maintenance load `delta Q_t` shall appear in routine reporting as levels, not as incident counts or period-over-period differences (Chapter 2, S2.9; Star & Ruhleder 1996).

S3.13 — Scheduled expiry for the high-`delta` portfolio. Components with target `delta >= 0.20` shall carry an explicit expiry date, after which continuation requires an affirmative decision. Silence retires them.

S3.14 — Escalate model breakdown. A fitted `a = 1−delta` outside `[0,1]`, or a decay window whose residuals are structured rather than random, shall be escalated as a specification failure and re-modelled. It shall not be clipped or smoothed.

The Failure Modes

THE PERSON-SHAPED REPOSITORY — a substantial fraction of stored order lives in one individual and nowhere else; `phi` is near 1 for that component.

- *Parameter damaged:* `delta` via the turnover channel; `C` via onboarding.

- *Diagnostic signature:* a name appears as the answer to "who would know?" for many unrelated questions. Their absence stops work in areas unrelated to their nominal role. They cannot take leave.
- *Measurement:* sample the "who would you ask about X" network across a fixed question set and compute the concentration. Compute `delta_turnover = p * phi` for the affected component.
- *Fix:* train to three holders (S3.11), then move the content down the artifact ladder. Note that the same actions raise `eta` — Chapter 4's mechanisms and this chapter's are the same mechanisms.

THE ROTTING DOCUMENT — the artifact is intact, complete, well-formatted, and false. Nothing about its appearance indicates this.

- *Parameter damaged:* `delta` via documentation entropy and environment drift, both undetected.
- *Diagnostic signature:* documents are never edited and never deleted. When someone follows one, it fails at step four, and the response is to ask a person rather than to fix the document.
- *Measurement:* select documents at random and have someone who did not write them execute them. Record the failure rate and the step at which failure occurs. That rate, per period, is a direct `delta` estimate for the documentation channel.
- *Fix:* make it executable (S3.10). Where that is impossible, date it and expire it (S3.13): a document with a visible staleness marker is honest about its own `delta`.

THE REORG AMNESIA — a structural change invalidates the part of `Q` encoded in the communication topology, and nobody accounts for it because no document was lost.

- *Parameter damaged:* `delta`, in a single large step rather than a rate.
- *Diagnostic signature:* post-reorganisation, routine requests take substantially longer and get routed wrongly. Everyone's documents are intact. Conway's Law (1967) is the mechanism.
- *Measurement:* the before/after network sample of S3.7. The fraction of changed or lost answers is the estimate.
- *Fix:* budget the loss in advance and pay it deliberately — publish the new routing, run the "who would you ask" survey and publish its answers as an artifact, and expect a recovery period of roughly `t_95` at the prevailing `delta`. Do not model the step as a rate; §3.7's cliff pitfall applies.

THE IMMORTAL RULE — a procedure whose rationale has decayed while the procedure has not. It cannot be retired because nobody can establish that retiring it is safe.

- *Parameter damaged:* `delta` far too low on a component that should have been high; `C` permanently elevated.
- *Diagnostic signature:* asked why the rule exists, everyone gives a different answer or cites a person who has left. The rule survives every review because removing it requires proving a negative.

- *Measurement:* for each rule in the active set, record whether any current member can state its originating condition. The fraction that cannot is the context-loss share of Q, and it is usually startling.
- *Fix:* sunset clauses applied retroactively — set an expiry date on the existing rule set and require affirmative renewal (S3.13). This converts "prove it is safe to remove" into "state why it should stay," which is the tractable direction.

THE PRESERVED GRUDGE — an interpersonal record maintained at near-zero `delta` in a component whose target `delta` should be the highest in the portfolio.

- *Parameter damaged:* `delta` too low; u driven toward 0 within the affected pair, because the interaction becomes fully predictable and defensive; `C` elevated by routing work around the pair.
- *Diagnostic signature:* two people or two groups who are never scheduled together, whose work is routed through an intermediary, and whose original incident predates most of the current membership.
- *Measurement:* count the routing detours. Compare the age of the originating incident against the half-life you would assign this component in the portfolio (§3.13 suggests 1.4 months).
- *Fix:* the design lever is the `delta` of the *record*, not the sentiment. Do not retain durable written accounts of interpersonal incidents past their intended lifetime. Where the record has already outlived it, a structured repair with a stated closure is a one-off high-`delta` event. **CANON §5.4 applies without exception: no mechanism here removes anyone's ability to decline.**

THE INVISIBLE MAINTENANCE — the activity that suppresses `delta` produces non-events, is therefore invisible to difference-based reporting, and is therefore chronically underfunded.

- *Parameter damaged:* `delta` rises as maintenance is cut; the rise is not attributed to the cut because the lag is roughly `t_half`.
- *Diagnostic signature:* maintenance budget falls with no immediate consequence, followed some months later by a cluster of failures treated as unrelated incidents. Star & Ruhleder (1996).
- *Measurement:* report `delta` and `delta Q_t` as levels (S3.12). Correlate maintenance spend against `delta` with a lag of one half-life.
- *Fix:* instrument the non-event. Scheduled drills, executable checks, and level-based reporting all convert invisible success into visible signal.

THE CONTAMINATED DECAY WINDOW — `delta` is estimated over a window in which input was still arriving, producing an underestimate.

- *Parameter damaged:* the estimate, and therefore every quantity computed from it — `Q_inf` and the Theorem 3.3 ceiling both have `delta` in the denominator.
- *Diagnostic signature:* the fit shows a plateau or an upturn; `R^2` is poor with structured residuals; per-step estimates trend toward zero across the window.

- *Measurement:* plot `ln Y_t` against `t` and inspect residuals for structure. Compare the regression estimate against the endpoint two-point estimate; disagreement above roughly 15% indicates non-geometric shape.
- *Fix:* find a cleaner window, shorten the window to its uncontaminated prefix, or fall back to the joint fit of Chapter 25 with its four required baselines.

THE ZERO-DELTA COMPONENT — a component of `Q` assigned no loss at all, deliberately or by neglect.

- *Parameter damaged:* `delta = 0` on that component: no fixed point, no error decay, unbounded linear accumulation (Theorem 3.3, second case).
- *Diagnostic signature:* the component's size has never decreased in its recorded history. Rule count, policy volume, or documentation size is monotone increasing. Onboarding time grows year over year with no corresponding growth in capability.
- *Measurement:* plot cumulative component size against time and look for any decrease. Compute the retirement rate; it is usually exactly zero.
- *Fix:* assign a target `delta` and an intended lifetime (S3.8). Institute scheduled expiry (S3.13). Accept that some structure will be lost that would have been useful — that loss is the price of the error decay and the fixed point, and Chapter 1 §1.7 computes what they are worth.

Worked Example

A twenty-two-person infrastructure team measures and re-engineers its `delta`

The system. Twenty-two engineers responsible for the operational infrastructure of a mid-sized service. Period: one month.

Step 1 — `Q` and its proxy `Y`. `Q` is operational stored order: the team's capacity to execute correct recovery, deployment, and diagnostic procedures without the original author present. The proxy `Y` is **the number of runbooks that pass a monthly spot-audit, where the audit consists of a randomly selected engineer executing the runbook end to end in a staging environment, and a pass requires completion without consulting the author.**

Proxy justification, per S3.5 and Paper E §6.3. The input observable `U` is the disagreement geometry across the team's seven operational categories in a given month — measured by error-decorrelation on incident post-mortems. The order observable `Y` is transferable executable capability, measured in a different environment, by a different person, at a later time, against a fixed procedure. `Y` cannot be raised by having more discussion; it can only be raised by producing or repairing a runbook and training someone to run it. **`Y` is not a repackaging of `U`.** This paragraph is in the design record before any fitting is performed.

Step 2 — The zero-input window. A six-month platform freeze during a compliance migration. During the freeze no new runbook work was funded, and the audit continued. This is a genuine zero-input window on the runbook component of `Q`.

month	0	1	2	3	4	5	6
Y (runbooks passing audit)	100	92	84	79	71	66	60

Step 3 — The fit. Log-linear regression of `ln Y` on `t`:

```
slope    = −0.08446
delta    = 1 − exp(−0.08446)  = 0.0810 per month
Y_0 fit  = 100.16
R^2      = 0.9980
t_half   = ln 2 / (−ln 0.9190) = 8.21 months
```

Cross-check by endpoints: `1 − (60/100)^{1/6} = 0.0816`, agreeing with the regression to 0.8%. The per-step estimates range from 0.0595 to 0.1013 — a reminder that no single step is an estimate. Residuals are unstructured and R^2 is 0.998, so the exponential form is supported for this proxy over this range.

In plain language for the sponsor: half of everything this team knows how to do that is not actively exercised is gone in eight months.

Step 4 — Channel decomposition, per S3.3. Monthly attrition is 1.4% (an annual rate of 15.6%), and an audit of runbook authorship and the "who would you ask" network puts the person-borne fraction at 0.55.

channel	mechanism	`delta_i`
turnover	`p * phi = 0.014 * 0.55`	0.0077
platform / tool churn	dependencies deprecate under the team	0.0200
environment drift	runbook references infrastructure that changed	0.0350
procedural forgetting	steps not executed between incidents	0.0200
naive sum		**0.0827**
combined `1 − prod(1−delta_i)`		**0.0804**
measured (regression)		**0.0810**
residual unexplained		**0.0006**

The residual is 0.0006, or 0.7% of the measured total — well inside S3.3's 25% threshold. **The enumeration is complete enough to act on**, and the largest channel is environment drift at 0.0350, 43% of the total.

Step 5 — The rest of the parameters. For completeness, since `delta` alone does not give a fixed point:

parameter	value	source
`N`	7 operational categories	enumeration (Chapter 7)
`u*`	0.45	error-decorrelation across post-mortems (Chapter 9)
`eta`	9.0 runbooks per nat per month	attributable `delta_Y` per unit S over 14 months (Chapter 4)

parameter	value	source
C	4.0 runbooks per month	deprecations, audit failures, forced rewrites (Chapter 5)

```
N_eff    = 1 + 6 * 0.45      = 3.70
S        = ln 3.70           = 1.3083 nats      (67.2% of ln 7 = 1.9459)
eta S    = 9.0 * 1.3083      = 11.7750 runbooks/month
eta ln N = 9.0 * 1.9459     = 17.5132 runbooks/month   (the per-month ceiling)
```

Step 6 — Model consistency check. Theorem 3.5 predicts

```
Y_inf = (eta S - C)/delta = (11.7750 - 4.0)/0.0810 = 96.00 runbooks
```

The observed pre-freeze steady state was 100. **The model predicts 96.00 against an observed 100 — a 4.0% discrepancy, which is well within the uncertainty of the parameter estimates and is reported as a consistency check, not as a validation.** Validation requires the frozen held-out procedure of Chapter 25, which is Step 10.

Step 7 — The intervention. Attack the largest channel first (§3.6), and use the artifact ladder (§3.11).

channel	before	intervention	after
turnover	0.0077	pair rotation on-call; three holders per runbook (S3.11) drops `phi` from 0.55 to 0.30	0.0042
platform / tool churn	0.0200	no change proposed — outside the team's control this cycle	0.0200
environment drift	0.0350	**runbooks executed monthly in CI against staging: prose becomes an executable check that fails when it becomes false (S3.10)**	0.0050
procedural forgetting	0.0200	quarterly exercise; each engineer executes two unfamiliar runbooks live	0.0080
combined	**0.0804**		**0.0368**

```
new delta   = 0.0368 per month
new t_half  = ln 2 / (-ln 0.9632) = 18.49 months     (was 8.21)
```

Step 8 — What that buys, and what it costs, per S3.9.

```
Y_inf before = (11.7750 - 4.0)/0.0810 = 96.00 runbooks
Y_inf after  = (11.7750 - 4.0)/0.0368 = 211.28 runbooks     (+120.1%)

t_95 before  = ln(0.05)/ln(0.9190) = 35.5 months
t_95 after   = ln(0.05)/ln(0.9632) = 79.9 months            (+125.1%)
```

The team more than doubles its sustainable stock of executable capability, and more than doubles the time it takes to converge on any new steady state. Both numbers go in the proposal. The second one is the one that will not be volunteered by anyone else in the room.

Is that trade acceptable here? The relevant comparison is `t_95` against the rate at which the underlying platform changes. If the platform's own generation is shorter than 80 months, the team will be converging toward a target that has already moved, and the correct design is a lower target `Y_inf` with a faster response. In this case the platform generation is estimated at roughly 5 to 7 years, so 80 months is at the edge of acceptable and the design record says so explicitly rather than hiding it.

The added `C`: CI execution of runbooks costs compute and maintenance of the staging environment; quarterly exercises cost roughly 2 engineer-days per person per year; pair rotation costs scheduling flexibility. Estimated at 1.2 runbooks-equivalent per month of additional `C`. Re-running with `C = 5.2`:

```
Y_inf after, with added cost = (11.7750 - 5.2)/0.0368 = 178.67 runbooks  (+86.1%)
```

Still a large gain, and **the honest number is 178.67, not 211.28.** Proposals that omit the added `C` overstate the benefit by 18%.

Step 9 — Headroom, per Chapter 6. With `eta ln N = 17.5132` runbooks/month:

```
h = eta ln N - delta Y - C
```

state	delta	Y	C	h
before	0.0810	100	4.0	5.4132
after	0.0368	100	5.2	8.6332

and the stock at which headroom reaches zero, `Y = (eta ln N - C)/delta`:

```
before: (17.5132 - 4.0)/0.0810   = 166.83 runbooks
after:  (17.5132 - 5.2)/0.0368   = 334.60 runbooks
```

Before the intervention, the maintenance load `delta Y` at `Y = 100` is 8.10 runbooks per month — more than double the explicit cost `C = 4.0`. The team's largest recurring expense was invisible because it was not billed to anything. After the intervention it is 3.68. That single line is usually the one that carries the proposal, and it is Chapter 6's thesis arriving early: the maintenance term grows with what you have, and it is the term nobody puts on a budget.

Step 10 — Preregistration and falsification, per Chapter 25.

Fit `Y_{t+1} = a Y_t + b F(U_t) - C_hat_t + eps_t` with `0 <= a <= 1` and `b >= 0` on the 14 months preceding the intervention. Freeze. Predict the 24 months following.

Predicted: fitted a rises from `1 - 0.0810 = 0.9190` to `1 - 0.0368 = 0.9632`; b unchanged; Y rises toward 178.67 with 95% of the movement complete at month 80.

Falsified if:

1. fitted a after the intervention is not significantly greater than 0.9190 — the interventions did not reduce loss;
2. fitted a falls outside `[0,1]` — the first-order form is wrong for this proxy and the whole model must be re-specified;
3. b changes materially — the interventions have confounded `delta` reduction with `eta` increase, which is likely on its face, since executable runbooks and three-holder training plausibly raise conversion too. **This is a real confound and it is stated in advance rather than discovered afterward.** The mitigation is to stage the interventions: CI execution in quarter one, three-holder training in quarter three, with a fit between them.
4. the frozen model fails to beat the **shuffled-input control** on held-out data — which would mean the timing of interaction carries no information beyond its volume, and the storage-law description of this system is not supported regardless of how well `delta` was measured.

The shuffled-input control is the one most often omitted and the one most often decisive. It is the only baseline that can distinguish a real dynamical model from a well-fitted trend.

Step 11 — The `delta` portfolio for this team, per S3.8. The intervention above concerns one component. The full record assigns a target to each:

component	target `delta`/month	half-life	intended lifetime
incident-response protocol	0.010	69.0 months	multi-generation; safety-critical
runbooks	0.037	18.5 months	one platform generation
tooling conventions	0.050	13.5 months	one platform generation
on-call rota practice	0.100	6.6 months	re-decided twice yearly by exercise
quarterly roadmap	0.250	2.4 months	must go stale visibly
post-incident personal attributions	0.400	1.4 months	cleared by design; see THE PRESERVED GRUDGE

No row is zero (R3.3), every row has a stated lifetime (R3.4), and the two high-`delta` rows carry expiry dates so that silence retires them (S3.13).

What this chapter has established

`delta` is the fraction of stored order lost per period without maintenance. It is not observed net decline, it is not a constant of nature, and it is not a single quantity — it is a set of named channels that compose as `1 − prod(1−delta_i)`, measurable from a zero-input window by a log-linear fit, and separately engineerable channel by channel.

Driving it down means moving stored order down the artifact ladder toward media that fail loudly when they become false, and it carries two prices that must be stated together: `C` rises, and `t_95`

rises in exact proportion to the gain in `Q_inf`.

Driving it up is a legitimate and underused move. Grudges, obsolete procedure, stale plans, and captured metrics all need high `delta`, and the mechanism — scheduled expiry, so that silence retires — is the cheapest removal tool available. Every component needs a target `delta` and an intended lifetime, and no component gets zero, because at `delta = 0` there is no fixed point, no error decay, and no correction: the first column of §3.12's table is a column of ones.

A system with no forgetting cannot learn. Chapter 1 established that the leak is what makes arrival possible; this chapter establishes that the leak is also what makes *revision* possible, and that the two are the same property read at different timescales.

Chapter 4 turns to `eta` — how much of what happens becomes durable at all. It is the most neglected parameter in this book's experience, usually the cheapest to improve, and, as §3.14 computed, the lever that beats `delta` reduction in the general case by a margin of 32 percentage points.

Chapter 4 . Conversion — Engineering Eta

The Law

The storage law from Paper E, §3, reproduced without alteration:

```
Q_{t+1} = (1 − delta)∗Q_t + eta∗S_t − C_t      S_t = F(U_t) = ln(1 + (N−1)∗U_t)

Q_t   stored order  (institution, habit, furnished room, trust, working method)
S_t     surplus input this period,  0 <= S_t <= ln N
delta   retention loss in [0,1]
eta     conversion efficiency >= 0
C_t     cost / dissipation >= 0
```

Chapter 3 dealt with `delta`, the rate at which stored order leaks away when nobody maintains it. This chapter deals with `eta`, the coefficient that decides how much of what happened becomes something that lasts.

`eta` is the only parameter in the law that stands between an interaction and a consequence. Everything else describes what happens to consequences once they exist. `eta` decides whether they exist at all.

4.0.1 What is proven and what is postulated

Badge the pieces plainly, because the rest of the chapter leans on the distinction.

Proven (Paper A, Theorem 2.1): `F(x,y) = ln(1 + (N−1)u)` with `u = 1 − (x.y)^2`, uniquely characterised by axioms S1–S4. **Proven (Paper D, Corollary 2.2):** `0 <= F(u) <= ln N`. **Proven (Paper E, Theorem 3.1, 3.3, 3.5; Corollaries 3.2, 3.4; Proposition 4.1):** the exact solution, the uniform bound, the steady state, the one-step growth condition, the sustainability ceiling, and the target-increment condition. These are mathematics. They do not depend on anything in this chapter being true about the world.

Postulated (design axiom S3, and stated as such in CANON §1.1): the affine effective count `g(u) = 1 + (N−1)u`. S3 is *chosen*. It is the axiom that makes the functional logarithmic rather than something else, and it is chosen because an affine effective count is the honest reading of "audit the headcount for redundancy before you count it."

Postulated (this chapter, badged): that a real organisation's conversion of interaction into durable structure is well modelled by a single scalar `eta` that is roughly constant over the horizon of interest. This is a modelling assumption, not a theorem. It is the assumption that makes the whole manual operational, and it is falsifiable — Chapter 25 gives the procedure. **If `eta` is not approximately stable across periods, the linear storage law is the wrong model for your system and you should say so rather than fit it anyway.**

Honesty clauses, repeated because the law is being used: the triangle inequality fails for geometric triples (Paper A §5); subadditivity holds on the parameter u only. As deployed, the surplus functional is a **lower bound** on the ensemble's information, not the information itself — we log the floor, not the ceiling. Paper A is pure mathematics and makes no empirical or metaphysical claim. Paper D does not identify surplus with Shannon entropy or mutual information. Paper E proves no thermodynamic negentropy law.

4.0.2 Units and range

`eta` has units of `[Q] / [S]`. `S` is dimensionless (a natural logarithm), so `eta` carries whatever unit the order proxy carries. In practice you will normalise `Q` so that `eta` lands somewhere in `[0, 1]`, but nothing in the law requires `eta <= 1`. `eta > 1` simply means your `Q` units are small relative to a nat of surplus. **Do not read `eta > 1` as a violation; read it as a unit choice.** What you may not do is compare `eta` values across systems that normalised `Q` differently. `eta` is only comparable within a fixed measurement convention (Chapter 24, §24.3).

4.0.3 Where eta appears

`eta` appears in every design inequality in the manual, and it always appears multiplied by something you cannot easily change:

```
one-step growth   (Cor 3.2):   Q_{t+1} > Q_t   <=>   eta*S_t > delta*Q_t + C_t
uniform bound     (Thm 3.3):   Q_t <= max{ Q_0 , eta*ln N / delta }
sustainability     (Cor 3.4):    nondecreasing requires  C_t <= eta*ln N -
delta*Q_t
steady state      (Thm 3.5):   Q_inf = (eta*F(u*) - c*) / delta
headroom          (§6.4):      h_t = eta*ln N - delta*Q_t - C_t
max increment     (Cor 4.2):   r_max = eta*ln N - delta*Q_t - C_t
```

Every one of these is linear in `eta` and logarithmic in `N`. That asymmetry is the whole argument of this chapter, and §4.9 prices it exactly.

4.0.4 The eta = 0 collapse

Set `eta = 0` in Theorem 3.1:

```
Q_t = (1-delta)^t * Q_0  -  sum_{k=0}^{t-1} (1-delta)^{t-1-k} * C_k
```

The driven term vanishes entirely. What remains is the decaying transient plus an accumulated debit. For `delta > 0` and any `C_k >= 0`, `Q_t` is non-increasing in every period and converges to `-c/delta` if cost is constant — which is inadmissible, meaning the linear model has been pushed outside its range and the real system has hit some floor (bankruptcy, dissolution, a hard reset).

A system with `eta = 0` does not stagnate. It declines, and it declines at exactly the rate it would decline if the interaction had never occurred. The meetings still happen. The conferences still convene. The retreat still costs three days. The arithmetic does not notice any of it.

This is the single most common condition in organisational life, and it is nearly invisible from the inside, because `S_t` is large and visible while `eta*S_t` is small and abstract. People experience the interaction. They do not experience the coefficient.

The Engineering

4.1 The distinction that the chapter turns on

There are two different quantities and ordinary language uses one word for both.

```
S_t      the interaction that HAPPENED   -- attendance, discussion, encounter, contact
eta*S_t   the interaction that STUCK     -- decision, artifact, default, trained person
```

S_t is bounded above by ln N (Paper D, Cor 2.2) and is generated by the geometry of who is in the room and how different they actually are. Chapters 7 through 11 engineer that geometry. eta is generated by what the room does with what it produced. **The geometry decides how much is available. The conversion decides how much you keep.**

Almost every organisation measures S and reports it as if it were eta*S. Attendance figures, meeting counts, conference days, "touchpoints," "engagement," workshop hours — these are all measurements of S, or of proxies for S. None of them is a measurement of Q. A report that lists them and concludes that the organisation is building something has performed a substitution the arithmetic does not permit.

The diagnostic phrase is in the manifest brief and it deserves to be stated as a rule:

"We had a great conversation" is a report of S. It contains no information about eta.

The correct follow-up is never "what did you discuss." It is: *what is different now that was not different before the conversation, and where is it written down.* If the honest answer is "nothing yet," then eta*S = 0 for that interaction, and the conversation, however good, moved Q by exactly zero.

4.2 What eta physically is

eta is a capture rate. In every real system it decomposes into three sequential gates, each of which is a multiplier in [0,1]:

```
eta  =  eta_record  *  eta_bind  *  eta_reach
```

- **eta_record** — the fraction of what was produced that was *recorded at all*, in any durable medium. A decision reached and not written is lost at the rate of human memory, which is a delta term with a half-life measured in days.
- **eta_bind** — the fraction of what was recorded that was *bound to a change in the system's behaviour*: an owner, a date, a default, a modified procedure, a moved wall. A record with no binding is an artifact in a filing cabinet. It has Q > 0 in the narrowest sense and no effect on anything.
- **eta_reach** — the fraction of the bound change that actually *reached the people and places that act*. A changed default that nobody's tooling implements has not changed anything.

Badge: this three-gate decomposition is a modelling convenience, not a theorem. It is useful because it tells you where to look. The multiplicative form is the important part: three gates at 0.7 each give eta = 0.343, and the intuition that "we're doing pretty well on all three" is worth less than half of what it feels like.

The three gates also explain why eta improvements are so often cheap. If your worst gate is eta_record = 0.1 because nobody writes anything down, then a scribe rota costs one person's attention per meeting and can plausibly take that gate to 0.8 — an eightfold move on the whole coefficient, for no change to headcount, geometry, or budget.

4.3 The six capture mechanisms

These are the durable-structure channels available to almost every system. They are listed in rough ascending order of `eta` per unit of effort, and each is described in the terms of the law it moves.

1. Writing it down (raises `eta_record`). The minimum viable capture. A shared, dated, append-oriented record of what was concluded. Cost: one person, ten minutes. **The single highest-leverage intervention available in any organisation with `eta_record < 0.5`, which is most of them.** Note the failure variant in §4.13: minutes that record *what was said* rather than *what was decided* raise `eta_record` while leaving `eta_bind` at zero.

2. The decision log (raises `eta_record` and `eta_bind` together). Distinct from minutes. A decision log records: the decision, the date, the alternatives rejected, the person accountable, the review date, and the condition that would reverse it. It is the cheapest device known that touches both gates at once. It also serves the falsification requirement of Chapter 25, because a decision log with reversal conditions is a preregistration.

3. The changed default (raises `eta_bind` and `eta_reach`). A default is a decision that executes itself. The template that now has the new field. The form that now asks the question. The calendar that now has the slot. **Defaults have the highest `eta_reach` of any mechanism, because they require no ongoing compliance from anyone.** This is Weiser's principle (1991) — "the most profound technologies are those that disappear" — applied to governance rather than computing. It is also Star & Ruhleder (1996): infrastructure is invisible when it works, which is exactly why a default is more durable than a policy.

4. The built object (raises all three, high `delta` resistance). Physical or software artifact: the jig, the checklist card, the script, the fixture, the moved partition. A built object is stored order that resists `delta` mechanically rather than through anybody's diligence. Chapters 12 through 17 are, from this chapter's point of view, an extended treatment of the built object as an `eta` device.

5. The trained person (highest `eta`, highest `delta` risk). Skill in a person is the densest form of stored order available. It is also the form most exposed to turnover, which Chapter 3 treats as a `delta` term. The correct engineering response is not to avoid training people but to pair mechanism 5 with mechanisms 1–4 so that the departure of a person is a partial rather than total loss.

6. The altered floor plan / altered institution (highest durability, highest cost, longest lag). Moving a wall, changing a reporting line, creating a standing role, rewriting a charter. These are the slowest and most expensive captures and they have the lowest `delta`. Conway's Law (1967) is the reason they work: systems mirror the communication structure of the organisation that builds them, so an alteration to the communication structure propagates into everything that structure subsequently produces.

4.4 Ordering the mechanisms honestly

The list above is ordered by *effort-adjusted leverage in a typical system*, and that ordering is a judgement, not a measurement. **The honest statement is that the ordering is system-specific and you should measure it rather than assume it.** What is not system-specific is the structure of the argument: find your lowest gate, and fix that one.

A concrete illustration of why the multiplicative structure matters. Consider a group whose gates are:

```
eta_record = 0.20   (someone occasionally emails notes)
eta_bind   = 0.60   (when notes exist, an owner is usually named)
eta_reach  = 0.70   (named owners usually follow through)

eta = 0.20 * 0.60 * 0.70 = 0.084
```

Now consider two candidate interventions of comparable cost:

```
Intervention A: improve eta_bind 0.60 -> 0.90
  eta = 0.20 * 0.90 * 0.70 = 0.126      (+50%)

Intervention B: improve eta_record 0.20 -> 0.80
  eta = 0.80 * 0.60 * 0.70 = 0.336      (+300%)
```

Both interventions are "process improvements." One is four times the other. The difference is entirely visible in advance if you decompose, and entirely invisible if you do not.

4.5 Ostrom's design principles read as eta-raising devices

Ostrom's *Governing the Commons* (1990) gives eight design principles for durable common-pool resource institutions, derived from field study of institutions that survived for centuries and institutions that failed. They are usually read as governance advice. Read as engineering, they are almost entirely `eta` and `delta` devices, and it is worth saying which is which.

Ostrom principle	Primary effect in the law	Mechanism
1. Clearly defined boundaries	Raises `eta_bind`; defines `N`	You cannot bind a decision to an undefined membership
2. Rules matched to local conditions	Raises `eta_reach`	A rule that does not fit is not executed
3. Collective-choice arrangements	Raises `eta` and protects `u`	Those affected participate; see Chapter 22
4. Monitoring	Raises `eta_reach`; lowers `delta`	Detects when stored order has decayed
5. Graduated sanctions	Lowers `delta`; bounds `C`	Maintains bindings without the cost of maximal enforcement
6. Conflict-resolution mechanisms	Lowers `C`	Conflict is a dissipation term (Chapter 5)
7. Recognition of rights to organise	Raises `eta_bind`	External non-recognition voids the binding
8. Nested enterprises (polycentric)	Raises `eta` at scale; bounds `C`	The topology answer of Chapter 5, §5.7

The most important reading is principle 4. Monitoring is not primarily a compliance device in Ostrom's data; it is the device that makes the state of `Q` *visible*. And a `Q` that is visible can be

maintained. Star & Ruhleder's observation that infrastructure becomes visible only on breakdown is precisely the failure that monitoring corrects, at a cost, and the cost is a `C` term. Principle 4 is therefore a deliberate trade: **spend `C` to lower `delta`.** That trade is only worth making when `delta*Q` exceeds the monitoring cost, which is a computable condition and is worked in Chapter 6.

Principle 8, nested enterprises, is the same structural answer that Chapter 5 reaches from the cost side and Chapter 8 reaches from Alexander's semilattice. Three independent lines arriving at the same topology is the strongest form of agreement this manual is able to offer, and it is still not a proof.

4.6 Apprenticeship

Apprenticeship is the highest-`eta` transfer mechanism in the human repertoire, and it is worth saying why in the terms of the law rather than by appeal to tradition.

An apprenticeship holds `S_t` at a moderate level continuously over a long horizon, with `eta_reach = 1` by construction — the person who receives the transfer is the person who will act. Writing has a reach problem (the reader may not be the actor); apprenticeship does not. Writing has a tacit-knowledge problem (what cannot be articulated is not recorded); apprenticeship does not, because the transfer channel is demonstration rather than description.

What apprenticeship has instead is a throughput problem. It is `O(1)` transfers per master per period, where a written artifact is `O(many)`. **This is the exact trade: apprenticeship buys a very high `eta` at a very low `S` throughput; documentation buys a much lower `eta` at very high throughput.** A system that uses only one of them is leaving the other's regime unserved. The craft trades that survived did both: a written standard *and* a term of years.

The design rule that follows is not "do apprenticeship." It is: **for the fraction of your stored order that is tacit, documentation has `eta_record` near zero regardless of effort, and no amount of writing will fix it.** Identify that fraction honestly and give it a demonstration channel.

4.7 What eta = 0 looks like from the inside

Named symptoms, so they can be recognised before the arithmetic is run.

- The same decision is re-litigated at intervals, and each re-litigation feels novel.
- Nobody can say when a standing practice was adopted or by whom.
- A new member's onboarding consists of a person telling them things.
- The phrase "we've talked about this" is used as an argument, unsuccessfully.
- Meeting outputs are described in terms of mood ("productive," "aligned," "good energy").
- There is a search problem: things are written down but cannot be found, which is `eta_record > 0` with `eta_reach ~ 0`.
- Departures cause disproportionate loss relative to the departing person's visible role.

The re-litigation symptom is the most diagnostic, because it is the direct observable signature of `Q` decaying back below the threshold where a decision counts as settled. It is also measurable: count decisions revisited per period, per Chapter 24's protocol.

4.8 Measuring eta

The definitional estimator, stated in the manifest brief:

```
eta_hat = delta_Q_attributable / S_input
```

Both terms are hard, and the honest procedure names both difficulties rather than hiding them.

Step 1 — choose the order observable `Y_t`. Paper E §5.1 requires five operational choices; the order observable is the third. `Y_t` is a measurable proxy for `Q_t`. Candidates: count of live written decisions with named owners; count of standing defaults implemented in tooling; number of procedures a randomly sampled member can execute unaided; count of built artifacts in active use.

Step 2 — justify that `Y_t` is not a repackaging of `U_t`. This is Paper E §6.3, and it is the most common measurement error in the whole framework. If your "order" measure is really a measure of how much interaction is happening, you have fitted the input to itself and your `eta` estimate is meaningless. Chapter 24 gives the full test. **The short form: if your order proxy would move on a week when a lot of talking happened and nothing was concluded, it is an interaction measure, not an order measure, and you must replace it.**

Step 3 — measure `S_input`. Estimate `N` (Chapter 7, effective-count audit) and `u` (Chapter 9, disagreement rate or error decorrelation), then compute `S = ln(1 + (N−1)u)` directly. Do not estimate S from attendance.

Step 4 — establish attribution. `delta_Q_attributable` is the part of the change in Y that would not have happened without the interaction. The minimum acceptable design is a comparison against the four baselines of Chapter 25: persistence, direct-input, matched autoregressive, and shuffled-input. **The shuffled-input control is the one that catches a false `eta`: if you shuffle which period's `S` is paired with which period's `delta_Y` and the fit does not degrade, you have measured a trend, not a conversion.**

Step 5 — fit the minimal model. Paper E §5.2:

```
Y_{t+1} = a*Y_t + b*F(U_t) − C_hat_t + eps_t     with 0 <= a <= 1, b >= 0
a = 1 − delta       b = eta
```

Fit on training data, **freeze**, and report held-out performance. A fitted `b` is your `eta`.

4.8.1 A worked estimation table Take a group with `N = 8` categories and `u = 0.30`, so `S = ln(1 + 7*0.30) = ln(3.10) = 1.1314` nats per period. The order proxy Y is normalised so that one unit equals one durable, owned, implemented decision still live at review. Observed attributable change per period, under six different capture regimes:

Capture regime	`delta_Q` per period	`eta_hat = delta_Q / 1.1314`
No artifact of any kind	0.000	0.000
Verbal agreement only	0.010	0.009
Notes emailed after the fact	0.050	0.044
Decision log with owner	0.140	0.124
Log + owner + changed default	0.280	0.247
Built artifact + trained person	0.450	0.398

Badge: these `delta_Q` figures are illustrative inputs chosen to span a plausible range, not measurements from a study. The arithmetic in the right column is real; the left column is what you must measure in your own system. This manual will not invent a citation for a number it has not measured. What the table demonstrates is the *shape* of the result — a factor of forty between the bottom and top regimes, produced entirely by what happens in the twenty minutes after the interaction ends.

4.9 The asymmetry — why eta is the first move

This is the quantitative core of the chapter. `eta` enters the budget linearly; `N` enters logarithmically; `u` enters through a concave saturating function. The consequence is a set of exact equivalences that are worth committing to memory.

4.9.1 Doubling eta is exactly squaring N The budget term is `eta * ln N`. Doubling `eta` gives `2*eta*ln N = eta*ln(N^2)`. Therefore:

To match a doubling of `eta` by adding categories, you must square `N`.

Current `N`	`N` required to match a doubling of `eta`
2	4
4	16
8	64
10	100
16	256
20	400
50	2,500
100	10,000

An organisation with eight genuinely distinct disciplines that wants to double its structural budget can either install a decision log, or find fifty-six more genuinely distinct disciplines. **The decision log is available this week.**

4.9.2 Doubling eta is often impossible to match by raising u To double `S = ln(1 + (N−1)u)` by raising `u` alone, you need `u'` such that `1 + (N−1)u' = (1 + (N−1)u)^2`, giving

```
u' = ( (1 + (N−1)u)^2 − 1 ) / (N−1)
```

Since `u <= 1` by construction (`u = 1 − (x.y)^2` and `u` is in `[0,1]`, Paper A), this is feasible only when `(1 + (N−1)u)^2 <= N`, that is:

```
u  <=  1 / (sqrt(N) + 1)
```

Above that threshold, doubling the surplus by raising non-parallelism is not merely expensive — it is arithmetically impossible, because there is not enough room left below `u = 1`.

N	maximum u from which doubling via u is still feasible
2	0.4142
3	0.3660
4	0.3333
5	0.3090
8	0.2612
10	0.2403
16	0.2000
20	0.1827
50	0.1239
100	0.0909

Worked instances:

```
N=8,  u=0.30 -> S = 1.1314;  doubling needs u' = 1.2300   INFEASIBLE (u' > 1)
N=8,  u=0.10 -> S = 0.5306;  doubling needs u' = 0.2700   feasible, and expensive
N=20, u=0.10 -> S = 1.0647;  doubling needs u' = 0.3900   feasible, and very expensive
```

Raising u from 0.10 to 0.27 in a real team means replacing or retraining a substantial fraction of it — recruitment across training lineages, rotation, adversarial roles, the deliberate outsider (Chapter 9). That is a multi-year programme with a high failure rate. Doubling `eta` from 0.084 to 0.168 by fixing the recording gate is a rota.

4.9.3 The full comparison, one base case Base: `N = 8`, `u = 0.30`, `eta = 0.10`, `delta = 0.10`, `C = 0.05`. Then `S = 1.1314` and the structural budget is `eta * ln N = 0.10 * 2.0794 = 0.2079`.

Move	new `eta*ln N`	gain	realistic cost
baseline	0.2079	—	—
`eta`: 0.10 -> 0.20	0.4159	+0.2079	a scribe rota and a decision log
N: 8 -> 10	0.2303	+0.0223	two genuinely new disciplines hired
N: 8 -> 16	0.2773	+0.0693	doubling the discipline count
N: 8 -> 64	0.4159	+0.2079	eight times the discipline count

Hiring eight times as many distinct disciplines buys exactly what a decision log buys. This is the asymmetry, stated as a single line, and it is why `eta` is the first move in almost every real system.

4.9.4 Cadence is not an eta move, and must not be booked as one A fourth move gets proposed in the same breath as the three above — *meet more often* — and it belongs in a different column of the ledger. Chapter 59 settles the accounting and Chapter 2 §2.10 carries the correction:

for a linear time-invariant integrator, the integrated effect of one large input and of the same total delivered periodically is identical to every digit (433.3333 in all three regimes Chapter 59 tests). Rescheduling delivers no extra input, and a cadence change booked as a gain in `eta` is double-counting.

What cadence does buy is real, and it arrives through the *surplus* term, not the conversion term. Because `f(u) = ln(1 + (N−1)u)` is strictly concave with `f(0) = 0` (Paper A, Theorem 3.3; axiom S2), Theorem 59.1 gives that splitting a fixed budget across more occasions strictly increases the delivered total. At this chapter's base case `N = 8`, a session's worth of non-parallelism spread across `n` occasions:

occasions `n`	`u` per occasion	`f(u)`	delivered total `n f(u)`
1	0.60	1.6487	1.6487
2	0.30	1.1314	2.2628
4	0.15	0.7178	2.8714
6	0.10	0.5306	3.1838

Four occasions deliver **1.74 times** what one delivers, from the same budget — and separately, Paper D's ceiling `S <= ln N` clips whatever the single occasion tried to exceed (in Chapter 59's worked annual rite, **97.2% of the nominal budget is not delivered at all**). **Cadence beats intensity because of concavity and the ceiling, not because it delivers more input.** Book the gain under `S`, keep `eta` for the capture gate, and the two moves stay separately auditable.

4.10 Where eta cannot be raised

Three honest limits.

Limit 1 — the tacit fraction. Some stored order is not articulable. Documentation has `eta_record` near zero on it by nature. §4.6 gives the answer: a demonstration channel.

Limit 2 — the capture tax. Capture consumes attention, and attention is a `C` term. A system that captures everything spends its entire budget on capture. The marginal condition is computable: capture is worth doing on an interaction while the marginal `delta_eta * S` exceeds the marginal `delta_C`. In practice, this means **capture the decisions, not the discussion.** The discussion is `S`; only the decision can become `Q`.

Limit 3 — the ceiling is still `ln N`. Raising `eta` scales the whole budget, but the budget is still `eta * ln N`, and if `ln N` is small because `N` is small or `u` is near zero, then a high `eta` on nearly nothing is nearly nothing. **`eta` is a multiplier, and multipliers cannot rescue a zero.** Chapter 11's Parallel Wall — three glasses ground from one doctrine, `u = 0`, `N_eff = 1`, `S = ln(1) = 0` — is exactly the case where every `eta` intervention in this chapter returns nothing, because `eta * 0 = 0` for all `eta`.

This gives the correct sequencing rule, which Chapter 26 formalises:

Check that `S > 0` before investing in `eta`. If `u ~ 0`, fix the geometry first. Otherwise fix `eta` first, and fix it before you consider `N`.

4.11 Conway's Law and the eta of the channel

Conway (1967): a system's structure mirrors the communication structure of the organisation that produced it. In the terms of this law, Conway's observation is a statement about `eta_bind` and `eta_reach`: **the organisation's communication topology is the channel through which captured order reaches the artifact, so the artifact inherits the topology.**

Two engineering consequences.

First, if you want an artifact with a particular structure, the cheapest route is to give the producing organisation that structure — because `eta_reach` along the existing communication topology is high and `eta_reach` across it is low. Fighting Conway costs `C`; using it is free.

Second, and less often noticed: **Conway's Law implies that a communication topology change is a high-`eta` intervention**, because it propagates automatically into everything subsequently produced. Mechanism 6 in §4.3 (the altered institution) is expensive precisely because it is powerful, and Conway's Law is the reason it is powerful.

4.12 The dark side: high eta on the wrong thing

`eta` is neutral. It converts whatever surplus arrives into whatever durable structure the capture mechanism specifies. A high-`eta` system that captures the wrong conclusions builds durable wrong structure, and durable wrong structure has a low `delta`, which means it persists.

This is the mirror image of Chapter 3's observation that **a system with no forgetting cannot learn.** A system with very high `eta` and very low `delta` converges quickly to a fixed point and then cannot leave it. Banach guarantees the arrival (CANON §1.3); Banach says nothing about whether the fixed point is a good place.

The engineering answer is not to lower `eta`. It is to build reversal into the capture format itself. This is why the decision log specification in §4.3 includes *the condition that would reverse it*. **A capture format that records its own falsification condition converts surplus into structure that can be dismantled on evidence, which is the only kind of high-`eta` system that remains correctable.** Chapter 25 makes this a formal requirement; Chapter 6, §6.9 makes the dismantling itself a design lever.

Goodhart's Law and Campbell's Law belong here too. If `eta` is measured, and the measurement becomes a target, then capture ceremonies will multiply and the measured `eta` will rise while the real one does not. The defence is the order-proxy justification of §4.8 step 2, applied adversarially: **ask whether your `eta` metric would rise if people simply wrote more documents that changed nothing. If it would, it is not measuring `eta`.**

4.13 The distinction between minutes and a decision log

Worth its own subsection, because the confusion is nearly universal.

```
MINUTES                              DECISION LOG
records what was SAID                records what was DECIDED
chronological                        indexed by decision
complete                             selective
raises eta_record only               raises eta_record AND eta_bind
read once, by nobody                 read on every re-litigation
```

```
grows without bound                grows with decision count
no owner field                     owner field is mandatory
no reversal condition              reversal condition is mandatory
```

Minutes are an archive. A decision log is a control surface. **The test is whether the document is consulted when someone asks "why do we do it this way." If the answer is no, you have minutes, and your `eta_bind` is unaffected by them.**

The Specification

Numbered, checkable requirements. A system claiming to have engineered `eta` satisfies these.

C4.1 The system shall define an order observable `Y_t` in writing, together with its measurement procedure and its period length.

C4.2 The system shall document the argument that `Y_t` is not a repackaging of the interaction observable `U_t` (Paper E §6.3). This argument shall be specific, not asserted.

C4.3 Every recurring interaction of more than two participants shall have a named recorder before it begins. The role may rotate; it may not be vacant.

C4.4 The system shall maintain a decision log. Each entry shall contain, at minimum: decision, date, alternatives rejected, accountable person, review date, reversal condition.

C4.5 No entry in the decision log shall have a vacant accountable-person field. An entry without an owner is `eta_bind = 0` and shall be marked as a proposal, not a decision.

C4.6 Every decision log entry shall carry a reversal condition — the observation that would require the decision to be revisited. Entries without one shall not be closed.

C4.7 The system shall maintain a register of changed defaults, distinct from the decision log, listing for each: the default, the tooling or artifact that implements it, and the date it became live.

C4.8 A decision that has not been implemented in tooling, template, form, calendar, or physical arrangement within one review period shall be reported as unrealised. Its contribution to `Q` is zero until then.

C4.9 The system shall estimate `eta` at least once per twelve periods, using the estimator `eta_hat = delta_Q_attributable / S_input`, with `S` computed from measured `N` and `u`, not from attendance.

C4.10 Every `eta` estimate shall be reported alongside all four baselines of Chapter 25: persistence, direct-input, matched autoregressive, and shuffled-input.

C4.11 The system shall decompose `eta` into `eta_record`, `eta_bind`, `eta_reach` at each estimation, and shall direct the next intervention at the smallest of the three.

C4.12 The system shall identify, in writing, the fraction of its stored order it considers tacit, and shall provide a demonstration channel (apprenticeship, pairing, shadowing) for that fraction.

C4.13 No capture requirement shall be imposed whose estimated `C` exceeds its estimated `delta_eta * S`. Capture is subject to the same budget as everything else.

C4.14 Capture shall target decisions and changed defaults. Verbatim capture of discussion shall not be a requirement of any recurring interaction.

C4.15 The system shall count re-litigated decisions per period as a standing diagnostic. A rising count is the signature of `eta*S < delta*Q`.

C4.16 Onboarding shall be defined as a document plus a demonstration channel, never as a sequence of conversations.

C4.17 Before any `eta` programme is funded, the system shall verify `S > 0` by measuring `u`. If `u` is near zero, the `eta` programme shall be deferred until Chapter 9's geometry work is done.

C4.18 The system shall not report attendance, meeting count, workshop hours, or engagement as evidence of `Q`. These are `S` proxies and shall be labelled as such wherever they appear.

C4.19 Any `eta` metric adopted as a target shall be re-examined for Goodhart capture at each estimation, using the test in §4.12.

C4.20 Every capture mechanism shall be reversible by the people subject to it. No capture requirement shall remove a participant's ability to decline participation (CANON §5.4).

The Failure Modes

Each entry: name, clinical description, damaged parameter, diagnostic signature, detecting measurement, fix. This format is the one used by the full catalogue in Chapter 27.

FM-4.1 . The Leaky Vessel

Damages: `eta`. **Description.** Interaction of high quality and adequate geometry occurs on schedule, and none of it is captured. `S` is large; `eta` is near zero; `Q` decays at rate `delta` regardless. **Signature.** Recurring re-litigation of settled questions; institutional knowledge held entirely in specific individuals; onboarding by conversation; the phrase "we had a great conversation" used as an outcome report. **Detection.** Count decisions revisited per period. Ask five members independently to state the current policy on a specific matter and compare answers; disagreement is direct evidence of `eta_record` failure. **Fix.** Named recorder (C4.3), decision log (C4.4). Cheapest fix in the manual.

FM-4.2 . The Minutes Fallacy

Damages: `eta_bind`. **Description.** Extensive verbatim records exist. Nothing is bound. `eta_record` is high, `eta_bind` is near zero, and the product is near zero. **Signature.** A large, complete, chronological archive that nobody consults. Minutes approved at the following meeting as a ritual. No owner fields anywhere. **Detection.** Access logs on the archive. Ask when a minute was last cited in an argument about what to do. **Fix.** Replace minutes with a decision log (§4.13). Keep the minutes if you like them; they are not the control surface.

FM-4.3 . The Write-Only Archive

Damages: `eta_reach`. **Description.** Capture occurs and binding occurs, but the record cannot be found by the people who need it. Search failure, permission failure, or format failure. **Signature.** People re-derive things that are documented. New members are told "it's in the wiki somewhere." Multiple contradictory documents on the same topic, each authoritative-looking. **Detection.** Give a new member a specific question and time how long retrieval takes. Above five minutes, `eta_reach` is compromised; above twenty, it is zero in practice. **Fix.** Single index, mandatory supersession marking, one canonical location per topic. Deleting duplicates is a `Q`-shedding action and is legitimate (Chapter 6, §6.9).

FM-4.4 . The Unowned Artifact

Damages: `eta_bind`. **Description.** A decision or artifact exists with no accountable person. It degrades without anyone being responsible for noticing. **Signature.** "Someone should…" as the standing conclusion. Orphaned tooling that everyone depends on and nobody maintains. **Detection.** Audit the decision log and the tooling register for vacant owner fields. The count is the measurement. **Fix.** C4.5. An unowned entry is a proposal, and shall be labelled one.

FM-4.5 . The Capture Tax

Damages: `C`, indirectly `eta`. **Description.** Capture requirements proliferate until the cost of documenting exceeds the value of what is documented. Sessions produce three artifacts about the session and none about the work. **Signature.** Meetings about the meeting format. Templates with more fields than content. Rising `C` with flat `Q`. **Detection.** Time-audit: person-hours spent on capture as a fraction of person-hours spent in interaction. Above roughly a fifth, examine C4.13 seriously. **Fix.** C4.13 and C4.14. Capture decisions and defaults only. Delete fields nobody reads.

FM-4.6 . The Ceremony That Replaced the Capture

Damages: `eta`, while appearing to raise it. **Description.** A capture ritual is performed faithfully and produces no binding. Retrospectives that generate action items nobody executes; postmortems filed and closed. **Signature.** Action items with a completion rate below a third and no consequence for non-completion. Identical items recurring across sessions. **Detection.** Track action-item completion and recurrence rate. Recurrence is the tell. **Fix.** Attach every item to the changed-default register (C4.7) or delete it. An item that cannot be expressed as a changed default, a built object, or a trained person is not a capture.

FM-4.7 . The High-Eta Error

Damages: nothing in the law; the law is working correctly against you. **Description.** A system with high `eta` and low `delta` has converted a wrong conclusion into durable structure and cannot dislodge it. Banach guarantees arrival at the fixed point; nothing guarantees the fixed point is desirable. **Signature.** Confident, well-documented, deeply implemented practice that no evidence appears able to change. Challenges are met with the documentation. **Detection.** Check the decision log for reversal conditions. Their absence is the diagnostic. **Fix.** C4.6 retroactively. Add reversal conditions to existing entries. Where a reversal condition has already been met and the decision stands, that is a `Q`-shedding case for Chapter 6, §6.9.

FM-4.8 . The Attribution Illusion

Damages: the `eta` estimate, not `eta` itself. **Description.** `eta` is estimated by correlating interaction with subsequent order growth, with no control. Any secular trend produces a positive estimate. **Signature.** An `eta` estimate that has never been tested against a shuffled-input control. A model that fits everything and predicts nothing out of sample. **Detection.** Run the shuffled-input control (Chapter 25). If fit does not degrade, the estimate is spurious. **Fix.** C4.10. Report all four baselines or report no estimate.

FM-4.9 . The Order Proxy That Was Really the Input

Damages: every downstream number. **Description.** `Y_t` is a disguised measure of `U_t` — for instance, "collaboration index" or "cross-team touchpoints." The fitted `eta` measures a definition, not a conversion. **Signature.** Y rises on weeks with much discussion and no conclusions. **Detection.** Paper E §6.3's test, given in full in Chapter 24. Simplest version: identify a period of intense interaction that produced nothing durable, and check whether Y rose. **Fix.** C4.2. Choose a proxy that can fall while interaction is high.

FM-4.10 . Eta Applied to a Parallel Wall

Damages: the intervention budget. **Description.** A serious capture programme is installed in a group with u near zero. `S = 0`, so `eta*S = 0` for any `eta`, and the entire programme returns nothing while consuming `C`. **Signature.** Excellent documentation of unanimous conclusions. A decision log in which nothing was ever contested. **Detection.** Measure u (Chapter 9). Run the disagreement audit (Chapter 11). **Fix.** C4.17. Defer the `eta` programme; fix the geometry first.

Worked Example

A nine-person research group over twenty-six weeks

System. A nine-person applied research group. Weekly periods. The group holds a ninety-minute all-hands each week and a variable number of pair sessions.

Survey (Chapter 24 procedure, abbreviated).

Categories. Enumerating genuinely distinct methodological lineages rather than job titles, the group has `N = 6`: experimental measurement, numerical simulation, statistical inference, hardware fabrication, domain theory, and software infrastructure. Three of the nine people duplicate an existing category, which is why the headcount is nine and `N` is six. This is the effective-count audit of Chapter 7 applied before anything else.

Non-parallelism. Disagreement audit over eight weeks: on structured forecasting questions about experimental outcomes, independent pre-discussion estimates disagreed at a rate giving an estimated `u = 0.55`. This is a genuinely healthy figure and reflects deliberate cross-lineage recruitment.

Surplus.

```
S = ln(1 + (N−1)*u) = ln(1 + 5*0.55) = ln(3.75) = 1.3218 nats/week
```

```
ceiling: ln N = ln 6 = 1.7918
geometric shortfall: ln N − S = 0.4700
```

The group is capturing 74% of the surplus its own category count would permit at `u = 1`. That is good geometry.

Retention loss. Measured by Chapter 3's method — `Q` observed at two points across a five-week shutdown with no input — `delta = 0.05` per week. A half-life of `ln 2 / (−ln(1−0.05)) = 13.5` weeks.

Cost. Tallied per Chapter 5: `C = 0.12` per week in normalised order-units, comprising all-hands coordination, review overhead, and administrative compliance.

Order observable. `Y_t` = count of live, owned, implemented methodological decisions and built artifacts in active use, normalised so one unit is one such item. `Y_0 = 0.80`. The proxy passes the §4.8 step-2 test: it fell during a busy but inconclusive fortnight in week 3–4 of the survey, which an interaction measure could not have done.

Diagnosis. Decomposing `eta`:

```
eta_record = 0.15   -- an occasional email; no standing recorder
eta_bind   = 0.40   -- owners named sometimes, review dates never
eta_reach  = 0.50   -- what is bound often fails to reach the people who act

eta = 0.15 * 0.40 * 0.50 = 0.030
```

The binding constraint is `eta`, and it is the recording gate specifically.

Before the intervention. With `eta = 0.03`:

```
eta*S    = 0.03 * 1.3218 = 0.0397 per week
delta*Q0 = 0.05 * 0.80   = 0.0400
C        = 0.1200

h_0 = eta*ln N − delta*Q_0 − C = 0.03*1.7918 − 0.0400 − 0.1200 = −0.1062
```

`h_0 < 0`. Corollary 3.4 applies without qualification: no arrangement of these nine people, no change to who talks to whom, no improvement in the quality of the discussion, can prevent `Q` from declining. The group's geometry is excellent and it is irrelevant, because the conversion coefficient is throwing away 97% of what the geometry produces.

Trajectory, `Q_{t+1} = 0.95*Q_t + 0.0397 − 0.12`:

week	Q_t	h_t
0	0.8000	-0.1062
1	0.6797	-0.1002
2	0.5653	-0.0945
4	0.3535	-0.0839
8	-0.0101	-0.0657
13	-0.3714	-0.0477
20	-0.7441	-0.0290
26	-0.9727	-0.0176

Q crosses zero at week 8. The steady state is `Q_inf = (0.0397 − 0.12)/0.05 = −1.607`, which is negative — Theorem 3.5 states plainly that a negative equilibrium means the regime is unsustainable or the linear model must be replaced. **The theorem does not hide this failure mode, and neither should the survey report.** In the real group, Q cannot go negative; what happens instead is that the group hits a floor consisting of whatever is externally enforced (funder reporting requirements, equipment logbooks) and operates entirely on that.

Note the deceptive shape of the h_t column. Headroom is *improving* every week — from -0.106 to -0.018 — and it is improving because Q is collapsing, which reduces the `delta∗Q` maintenance draw. **A system in decline shows rising headroom. Headroom rising while Q falls is not recovery; it is the maintenance burden shrinking because there is less left to maintain.** Chapter 6 §6.3 treats this inversion in full. It is the single most misread signal on the gauge.

Two thresholds, computed before choosing the intervention.

```
eta needed for a non−negative steady state:  eta > C/S = 0.12/1.3218 = 0.0908
eta needed for h_0 > 0 at Q_0 = 0.80:         eta > (delta∗Q_0 + C)/ln N
                                          = (0.0400 + 0.1200)/1.7918 = 0.0893
```

The group needs `eta >= 0.091`, a factor of three, and nothing more. That is the entire brief. It is worth noticing how much cheaper this brief is than the ones that would have been proposed without the arithmetic: not "hire more people" (N is already 6 and adding a seventh distinct lineage buys `eta∗ln(7/6) = 0.0046` at `eta = 0.03`), not "improve collaboration" (`u = 0.55` is already strong), not "reduce meetings" (cutting C to zero still leaves `h = 0.0537 − 0.0400 = +0.0137`, technically positive but a tenth of what the eta fix delivers, and it would destroy the u the group depends on).

Intervention. Three changes, all cheap, all targeting the decomposition:

1. *Standing recorder rota* on the weekly all-hands, rotating alphabetically, ten minutes of one person's time. Targets `eta_record`: 0.15 -> 0.70.
2. *Decision log* replacing minutes, with mandatory owner, review date, and reversal condition (C4.4–C4.6). Targets `eta_bind`: 0.40 -> 0.75.
3. *Changed-default register* (C4.7): every methodological decision must name the template, script, or fixture that implements it within one week, or be marked unrealised (C4.8). Targets `eta_reach`: 0.50 -> 0.80.

Projected: `eta = 0.70 ∗ 0.75 ∗ 0.80 = 0.420`. That is optimistic. Take the conservative half of the gain and project `eta = 0.21`, which is still more than double the threshold.

After the intervention. With `eta = 0.21`:

```
eta∗S = 0.21 ∗ 1.3218 = 0.2776 per week
h_0   = 0.21∗1.7918 − 0.05∗0.80 − 0.12 = 0.3763 − 0.0400 − 0.1200 = +0.2163
Q_inf = (0.2776 − 0.12)/0.05 = +3.1514
```

week	Q_t	h_t
0	0.8000	+0.2163
1	0.9176	+0.2104
2	1.0293	+0.2048
4	1.2362	+0.1945

week	Q_t	h_t
8	1.5914	+0.1767
13	1.9443	+0.1591
20	2.3084	+0.1408
26	2.5317	+0.1297

Q rises from 0.80 to 2.53 in twenty-six weeks and converges on 3.15. The time constant is `1/delta = 20` weeks; the group is at 80% of its new steady state by week 26, consistent with Chapter 2's settling-time table.

Notice what did *not* change: `N`, `u`, `S`, `delta`, `C`, and the headcount. **The same nine people, having the same conversations, with the same disagreement rate, at the same cost, moved from terminal decline to a fourfold increase in stored order — because someone writes it down and someone owns it.**

Falsification criteria, registered before the intervention (Chapter 25). The prediction is `a = 0.95`, `b = 0.21` in the minimal model, fitted on weeks 1–16 and frozen, with weeks 17–26 held out. The intervention is judged a failure if any of the following holds:

1. Held-out prediction of Y is not better than the persistence baseline `Y_{t+1} = Y_t`.
2. The fitted b is below 0.09 — the sustainability threshold — meaning the capture programme did not move the coefficient enough to matter even if it moved it.
3. The shuffled-input control fits as well as the true pairing, indicating the observed rise in Y is a trend the intervention did not cause.
4. Re-litigation count (C4.15) does not fall, indicating `eta_reach` is still zero despite records existing — FM-4.3 rather than success.

Standing caution. Y here counts owned, implemented decisions. That is a quantity a group could inflate by making trivial decisions and implementing them, which is Goodhart capture (C4.19). The registered defence is that the review-date field forces every entry to be re-examined, and entries that no longer correspond to anything anyone does are struck from the count. **A metric with a scheduled removal path is substantially harder to inflate than one that only accumulates**, and the removal path is also the shedding mechanism that Chapter 6 will show to be a headroom lever in its own right.

Chapter 5 . Cost — The Dissipation Budget

The Law

```
Q_{t+1} = (1 − delta)∗Q_t + eta∗S_t − C_t      S_t = F(U_t) = ln(1 + (N−1)∗U_t)

C_t     cost / dissipation >= 0
```

`C_t` is the term that removes stored order in a way that has nothing to do with decay. `delta` is what leaks out through the walls. `C_t` is what is drawn off deliberately: the overhead of being an organisation at all, paid every period, whether or not anything is accomplished.

Three results govern this chapter, all proven, all from Paper E, none altered.

Admissibility. `C_t <= (1−delta)∗Q_t + eta∗S_t` guarantees `Q_{t+1} >= 0`. This is the condition under which the linear model remains inside its own range.

Sustainability ceiling (Corollary 3.4). A nondecreasing step requires

```
C_t <= eta∗ln N − delta∗Q_t
```

and if `C_t > eta∗ln N − delta∗Q_t`, NO choice of interaction geometry can prevent order from decreasing. This is the single most important sociological statement in the book. It is a theorem, not a postulate. It says that once cost passes a computable line, no reorganisation of who talks to whom, no improvement in the quality of the difference between people, no better meeting, no more inspiring leadership, and no larger `N` can arrest the decline. The only remaining moves are to cut `C`, raise `eta`, or reduce `Q` — and Chapter 6 makes the last of those explicit.

Steady state (Theorem 3.5). With constant `U = u∗` and `C = c∗` and `delta > 0`:

```
Q_inf = (eta∗F(u∗) − c∗) / delta
```

and if `eta∗F(u∗) < c∗` the equilibrium is negative, meaning the regime is unsustainable or the linear model must be replaced. **The theorem does not hide this failure mode.** Neither does this chapter.

5.0.1 What is proven, what is arithmetic, and what is postulated

Badge them separately, because this chapter mixes three kinds of claim.

Proven mathematics: the three results above, plus Theorem 3.1 (exact solution), Theorem 3.3 (uniform bound `Q_t <= max{Q_0, eta∗ln N / delta}`), Corollary 3.2 (one-step growth), and Paper D Corollary 2.2 (`0 <= F(u) <= ln N`).

Arithmetic, therefore true, therefore not a claim about the world: in a group of `n` members, the number of distinct unordered pairs is `n(n−1)/2`. Brooks (*The Mythical Man-Month*, 1975) uses this count to explain why adding people to a late project makes it later. **The count is a fact of combinatorics. Brooks' *interpretation* — that coordination burden scales with this count — is an empirical claim about software projects, and Brooks presented it as observation from practice, not as a proof.**

Postulated here, badged: that `C` is well modelled as proportional to the number of active communication paths, `C = kappa ∗ P`, with `kappa` a per-link cost per period. This is a modelling

assumption of this manual. It is chosen because it is the simplest form that reproduces the observed superlinearity of coordination cost, and because it makes the crossover computable. **It is falsifiable: measure `C` at several `n` in a comparable setting and check whether the fitted exponent on `n` is near 2. If it is near 1, your organisation has already solved its topology problem and §5.7 explains how.**

Also postulated, and the more dangerous assumption: that the complete graph is the relevant path count. §5.7 shows it usually is not, and that the difference between the complete graph and a good topology is the largest single cost lever available to any organisation above about a dozen people.

5.0.2 Admissibility and sustainability are different constraints

A frequent confusion. They are not the same line and they are not close to each other.

```
admissible:     C_t <=  (1-delta)*Q_t + eta*S_t     -- Q_{t+1} stays >= 0
sustainable:    C_t <=  eta*ln N - delta*Q_t        -- Q_{t+1} >= Q_t
```

Admissibility is a floor condition: the model does not produce nonsense. Sustainability is a growth condition: order does not decline. **Admissibility permits eating your capital; sustainability does not.** And critically, `Q` appears with opposite signs in the two: a large stock of stored order *raises* the admissible cost (you can burn it) and *lowers* the sustainable cost (it must be maintained).

Worked, with `delta = 0.08`, `eta = 0.22`, `N = 8`, `u = 0.60` so `S = ln(1+7*0.60) = 1.6487` and `eta*S = 0.3627`, `eta*ln N = 0.4575`:

`Q_t`	admissible `C_max`	sustainable `C_max`
0.0	0.3627	+0.4575
0.5	0.8227	+0.4175
1.0	1.2827	+0.3775
2.0	2.2027	+0.2975
4.0	4.0427	+0.1375
8.0	7.7227	-0.1825

The two columns move in opposite directions, and at `Q = 8` the sustainable ceiling has gone negative while the admissible ceiling stands at 7.72. A large institution can therefore run for a very long time in a state that is admissible and unsustainable — consuming accumulated order at a rate that looks affordable because there is so much of it. This is the structural picture of long institutional decline, and it is the subject of Chapter 6.

The Engineering

5.1 What C actually contains

`C` is not one thing. Tally it as a sum of named lines, because an aggregate figure cannot be attacked and a decomposed one can. The standard taxonomy for this manual:

C1 . Coordination. The direct cost of aligning people: meetings, status reporting, scheduling, handoffs, approval chains, the negotiation of who does what. This is Brooks' territory and the one that scales with path count.

C2 . Friction. Cost incurred because the system is arranged badly rather than because the work is hard: waiting for access, re-entering data, tool mismatch, permission bottlenecks, finding out who to ask. Friction is C that produces nothing at all, not even coordination.

C3 . Maintenance of relationship. The upkeep that keeps u from collapsing into hostility or N from fragmenting. Real, necessary, and frequently unbudgeted. **This is the C line most often cut first and it is often the one that should be cut last**, because cutting it raises `delta` on trust, which is stored order.

C4 . Conflict. Disputes that consume attention without resolving. Note carefully: conflict and disagreement are *not* the same thing. Disagreement is u, which is on the credit side of the ledger. Conflict is C. **A design that suppresses conflict by suppressing disagreement has paid for a C reduction with a u reduction, and since u enters through `ln(1+(N−1)u)` and C enters linearly, that trade is usually bad.** Chapter 14's acoustic work and Chapter 17's table geometry are both, in part, devices for having disagreement without conflict.

C5 . Rent and premises. Literal rent, and everything that behaves like it: licence fees, platform tolls, the fixed cost of existing.

C6 . Commute and transit. The cost of getting bodies into proximity. This is a C line whose reduction directly trades against S, because proximity is one of the cheapest ways to raise interaction rate. It is a genuine trade and not a free cut.

C7 . Emotional labour. The cost of managing one's presentation, absorbing others' states, and maintaining the social surface. Real, chronically unmeasured, and unevenly distributed. It is a C term because it consumes exactly the resource — attention — that `eta` also consumes.

C8 . Compliance and reporting. Externally imposed overhead. Often the least compressible line, and the one that grows with Q (§5.11 and Chapter 6, §6.7).

C9 . Context switching and interruption. Each switch costs recovery time. This line is usually invisible in a budget and can be the largest single item in knowledge work. Chapter 15 treats it as a threshold problem: **an abrupt transition between zones of different character raises C, and a graded one lowers it.**

The tally is the first deliverable of any survey. Chapter 24 gives the worksheet.

5.2 Brooks' arithmetic

In a group of n people, the number of distinct pairs is

`P(n) = n(n−1)/2`

This is exact. The interpretive step — that coordination cost tracks P — is Brooks' observation, and this manual adopts it as a badged postulate (§5.0.1) in the form `C1 = kappa * P(n)`.

The count itself, alongside the surplus ceiling for the most favourable possible case where every added person brings a genuinely new category (`N = n`):

n	n(n−1)/2	ln n	paths per nat
2	1	0.6931	1.44
3	3	1.0986	2.73
4	6	1.3863	4.33
5	10	1.6094	6.21
6	15	1.7918	8.37
7	21	1.9459	10.79
8	28	2.0794	13.47
9	36	2.1972	16.38
10	45	2.3026	19.54
12	66	2.4849	26.56
15	105	2.7081	38.77
20	190	2.9957	63.42
25	300	3.2189	93.20
30	435	3.4012	127.90
40	780	3.6889	211.45
50	1,225	3.9120	313.14
75	2,775	4.3175	642.73
100	4,950	4.6052	1,074.88
150	11,175	5.0106	2,230.26

Read the last column as the price of a nat of surplus ceiling, denominated in communication paths. At `n = 6` a nat costs 8.4 paths. At `n = 150` — Dunbar's much-cited and contested figure — the same nat costs 2,230 paths. The price has risen by a factor of 266 while the thing purchased has risen by a factor of 2.8.

5.3 The divergence, stated per added person

The marginal view is more brutal than the total view, and it is the one that should govern hiring decisions.

Adding one person to a group of `n` creates `n` new paths and raises the surplus ceiling by `ln(n+1) − ln n`:

n → n+1	new paths	gain in ln n	paths per nat of gain
3 -> 4	3	0.28768	10.4
6 -> 7	6	0.15415	38.9
10 -> 11	10	0.09531	104.9
20 -> 21	20	0.04879	409.9
50 -> 51	50	0.01980	2,524.9
100 -> 101	100	0.00995	10,049.9
150 -> 151	150	0.00664	22,574.9

The marginal price of surplus rises quadratically in `n`, because the numerator grows like `n` and the denominator falls like `1/n`. From the fourth member to the hundred-and-fifty-first, the price of an additional nat rises by a factor of two thousand.

This is the divergence the manifest brief names: **cost grows quadratically while surplus grows logarithmically.** It is not a metaphor and it is not an approximation. `n(n−1)/2` is exactly quadratic and `ln n` is exactly logarithmic, and any quadratic overtakes any logarithm, at a point that depends only on `eta` and `kappa`. That point is the subject of §5.4, and it is the most important calculation in this book for team sizing.

5.4 The crossover calculation

Question. For given conversion efficiency `eta` and per-link cost `kappa`, at what group size `n` does the coordination cost `C = kappa * n(n−1)/2` exceed the entire structural budget `eta * ln N`?

Setup. Take the most favourable possible case for growth: every person added brings a genuinely distinct category, so `N = n`. This is the best case and it is almost never true; §5.6 handles the realistic case, in which the crossover arrives earlier and the *optimum* arrives much earlier. Solve for the smallest integer `n >= 2` such that

```
kappa * n(n−1)/2  >  eta * ln n
```

Result. The crossover group size `n*`, computed exactly:

kappa \ eta	0.10	0.20	0.30	0.50	0.80	1.00
0.0005	39	58	73	97	125	142
0.0010	27	39	49	66	85	97
0.0020	18	27	33	45	58	66
0.0050	11	16	20	27	35	39
0.0100	7	11	13	18	23	27
0.0200	5	7	9	12	16	18
0.0500	3	4	5	7	9	11
0.1000	2	3	4	5	6	7

How to read this table. `kappa` is the cost per active communication path per period, in the same normalised order-units as `Q`. `eta` is the conversion efficiency of Chapter 4. The entry is the group size at which coordination alone consumes the entire structural budget — the size at which, even with zero maintenance draw (`Q = 0`) and zero other cost, `h_t` has gone negative and Corollary 3.4 forecloses every geometric remedy.

Three things the table says that are worth stating flatly.

First: the crossover is an upper bound on viable size under a complete-graph topology, and it is a low number. For plausible mid-range values — `eta = 0.30`, `kappa = 0.01` — the ceiling is thirteen people. Not thirteen as a rule of thumb; thirteen as an arithmetic consequence of two measured parameters.

Second: `eta` buys size, and it buys it slowly. Reading across the `kappa = 0.01` row, a tenfold increase in `eta` from 0.10 to 1.00 moves the crossover from 7 to 27 — a factor of roughly 3.9, not 10. This is because `eta` multiplies a logarithm while `kappa` multiplies a quadratic. Setting `kappa*n^2/2 ~ eta*ln n` and ignoring the slowly-varying logarithm gives `n ~ sqrt(eta/kappa)` up to a log factor, so **the crossover scales roughly as the square root of `eta`.** Doubling `eta` buys about 41% more headcount, not 100% more.

Third: kappa buys size on exactly the same square-root terms, and kappa is far easier to move. Reading down the `eta = 0.30` column, a hundredfold reduction in `kappa` from 0.10 to 0.001 moves the crossover from 4 to 49, a factor of 12. And `kappa` is reducible by topology (§5.7), by protocol, by asynchrony, and by co-location, none of which require changing anyone's capability. **The cheapest way to be allowed to be bigger is to make each link cheaper, not to make each person better.**

5.5 The crossover is not the optimum — and the optimum is much smaller

The crossover answers "when does cost eat the whole budget." That is the wrong question for sizing. The right question is "what `n` maximises net structural gain," and the answer is a substantially smaller number.

Maximise `net(n) = eta*ln n − kappa*n(n−1)/2`. Differentiating and setting to zero:

```
d/dn [ eta*ln n − kappa*(n^2 − n)/2 ]  =  eta/n − kappa*(2n − 1)/2  =  0

  =>   eta = kappa * n * (n − 1/2)   ~   kappa * n^2

  =>   n_opt  ~  sqrt( eta / kappa )
```

Design rule: the net-optimal group size under a complete-graph topology is approximately the square root of the ratio of conversion efficiency to per-link cost. Exact integer optima, computed by direct search, with the net value in parentheses:

kappa \ eta	0.10	0.20	0.30	0.50	0.80	1.00
0.0005	14 (+0.218)	20 (+0.504)	25 (+0.816)	32 (+1.485)	40 (+2.561)	45 (+3.312)
0.0010	10 (+0.185)	14 (+0.437)	18 (+0.714)	23 (+1.315)	29 (+2.288)	32 (+2.970)
0.0020	7 (+0.153)	10 (+0.371)	13 (+0.613)	16 (+1.146)	20 (+2.017)	23 (+2.629)
0.0050	5 (+0.111)	7 (+0.284)	8 (+0.484)	10 (+0.926)	13 (+1.662)	14 (+2.184)
0.0100	3 (+0.080)	5 (+0.222)	6 (+0.388)	7 (+0.763)	9 (+1.398)	10 (+1.853)
0.0200	3 (+0.050)	3 (+0.160)	4 (+0.296)	5 (+0.605)	7 (+1.137)	7 (+1.526)
0.0500	2 (+0.019)	2 (+0.089)	3 (+0.180)	3 (+0.399)	4 (+0.809)	5 (+1.109)
0.1000	2 (-0.031)	2 (+0.039)	2 (+0.108)	3 (+0.249)	3 (+0.579)	3 (+0.799)

The closed form tracks the search closely: at `eta = 0.30, kappa = 0.005`, `sqrt(0.30/0.005) = 7.7` against a searched optimum of 8; at `eta = 1.0, kappa = 0.001`, `sqrt(1000) = 31.6` against a searched optimum of 32.

Compare the two tables at a common cell. At `eta = 0.30, kappa = 0.010`: **the optimum is 6 and the crossover is 13.** A team of thirteen is not "at capacity"; it is at *twice* its net-optimal size and delivering less than a team of six. Between 6 and 13 the group is still nominally viable — h may be positive — but every additional member is reducing net structural output.

This gap between optimum and crossover is where most organisations live, and it is invisible without the calculation, because nothing breaks at the optimum. The failure is silent and gradual until the crossover, at which point it becomes sudden and is misdiagnosed as a leadership problem.

Note also the bottom-left cell: at `eta = 0.10, kappa = 0.10`, the optimum is `n = 2` and the net value is **negative**. There exist parameter regimes in which no group of any size does better than nobody working together, because the per-link cost exceeds what any amount of difference can pay for. That is not a rhetorical flourish; it is what the arithmetic returns.

5.6 The realistic case: N saturates while n keeps growing

The tables above assume `N = n` — every hire brings a genuinely distinct category. This is the most generous possible assumption and it is false in essentially every real organisation. Real groups exhaust their genuinely distinct categories quickly. Beyond that point, additional people add paths without adding `N`, and Chapter 7's effective-count audit will show `N_eff` flat while headcount climbs.

Model it as `N = min(n, D)`, where `D` is the number of genuinely distinct categories available. Take `eta = 0.30`, `kappa = 0.01`:

D	surplus ceiling eta∗ln D	crossover n∗	net-optimal n	net at optimum
3	0.3296	9	3	+0.2996
5	0.4828	11	5	+0.3828
8	0.6238	12	6	+0.3875
12	0.7455	13	6	+0.3875
20	0.8987	13	6	+0.3875

Look at the last three rows. Raising the available category count from 8 to 20 — a substantial and expensive achievement — moves the optimum not at all, and moves the crossover from 12 to 13. Once the quadratic cost term is binding, extra categories are irrelevant because you cannot afford the people who would carry them. Chapter 8 reaches the same conclusion from the surplus side and states it as: **you cannot escape a headroom problem by adding domains.** Here it appears from the cost side, and the two derivations are independent.

Also note that when `D` is small, the optimum sits exactly at `D` (rows 1 and 2). **When categories are scarce, hire exactly as many people as you have categories and stop.** The eleventh person on a five-category problem is pure `kappa`.

5.7 Topology — you do not have to buy the complete graph

Everything above assumed `P = n(n−1)/2`: every pair maintains an active link. That is a choice, not a law, and it is the most expensive available choice. The path count is set by the communication topology, and the topology is a design parameter.

Path counts under five wiring plans, priced at `kappa = 0.006`:

n	complete n(n−1)/2	ring n	star n−1	tree n−1	cells of 6 + liaisons
6	0.0900	0.0360	0.0300	0.0300	0.0900
12	0.3960	0.0720	0.0660	0.0660	0.2040
18	0.9180	0.1080	0.1020	0.1020	0.3420
24	1.6560	0.1440	0.1380	0.1380	0.5040
30	2.6100	0.1800	0.1740	0.1740	0.6900
48	6.7680	0.2880	0.2820	0.2820	1.3920

The "cells" column assumes complete graphs of six within each cell, plus liaison links between cells priced at four times a normal link — because a cross-boundary link genuinely is more expensive, requiring translation, context transfer, and schedule reconciliation.

At n = 48, the complete graph costs 6.768 and the cellular topology costs 1.392 — **a factor of 4.86, purchased entirely by deciding who is not obliged to keep a live channel with whom.**

But topology is not free in the surplus term, and the honest treatment says so. The star and the tree are cheap because they route everything through a hub, and a hub is precisely the structure Chapter 8 identifies as driving u toward zero: **hierarchy is the organisational form that homogenises, because every message passes through a filter with its own single viewpoint.** Barabási & Albert (1999) show that preferential attachment produces hubs spontaneously; Albert, Jeong & Barabási (2000) show such networks are robust to random failure and fragile to targeted hub attack. A star topology therefore buys a cheap C at the price of u and of resilience. That is a real trade and it must be priced, not assumed away.

The cellular arrangement is the one this manual recommends, and the recommendation converges from four independent directions:

- **Cost (this chapter):** cells cap the quadratic term at cell size, making total C linear in cell count rather than quadratic in headcount.
- **Ostrom's eighth principle:** nested enterprises. Field-observed in long-lived institutions.
- **Alexander, "A City is Not a Tree" (1965):** a semilattice — overlapping, non-exclusive groupings — carries more structure than a strict tree. Cells with overlapping membership are a semilattice; cells with a reporting hierarchy are a tree.
- **Centola & Macy (2007) and Centola (2010):** complex contagion needs multiple reinforcing contacts, so clustered topology outperforms random long ties for spreading a behaviour. Cells are clusters.

Four independent literatures arriving at the same topology is the strongest agreement this manual can offer, and it is still not a proof. State it as convergent evidence and design accordingly.

The concrete rule that follows:

```
Keep complete-graph obligation inside cells of roughly n_opt = sqrt(eta/kappa).
Connect cells by explicit liaison links, budgeted at a higher kappa.
Let cell membership OVERLAP -- that is what makes it a semilattice rather than a tree.
Never route all inter-cell traffic through a single hub.
```

§5.13 works this arithmetic through on a real case.

5.8 Queueing — why slack is a cost-control device, not waste

Every shared resource in an organisation is a queue: the approver, the reviewer, the one person who knows the deployment procedure, the single lab instrument, the meeting room. Queueing theory gives a result that is uncomfortable and exact.

For an M/M/1 queue with utilisation `rho = arrival rate / service rate`, the expected waiting time in queue, measured in units of service time, is

```
Wq / service_time  =  rho / (1 − rho)
```

rho	wait, in service times	multiple of the wait at rho = 0.5
0.10	0.111	0.11
0.30	0.429	0.43
0.50	1.000	1.00
0.60	1.500	1.50
0.70	2.333	2.33
0.80	4.000	4.00
0.85	5.667	5.67
0.90	9.000	9.00
0.95	19.000	19.00
0.98	49.000	49.00
0.99	99.000	99.00

Between 90% and 99% utilisation — an 11% gain in apparent efficiency — the wait multiplies by eleven. Little's Law (`L = lambda * W`) then says the queue length multiplies by the same factor, so work-in-progress balloons, which raises context-switching cost (line C9) and coordination cost (line C1) simultaneously.

Therefore: waiting is a `C` term, and driving utilisation toward 1 converts idle capacity — which costs nothing in the law — into waiting, which costs a great deal. Slack is not waste. Slack is the instrument that keeps the queueing term out of `C`.

The design consequence is a hard number. If waiting costs `kappa_delay` per unit of service time waited, the delay contribution to `C` is `kappa_delay * rho/(1−rho)` per item. At `kappa_delay = 0.02`:

```
rho = 0.50  ->  0.0200 per item
rho = 0.70  ->  0.0467
rho = 0.80  ->  0.0800
rho = 0.90  ->  0.1800
rho = 0.95  ->  0.3800
rho = 0.99  ->  1.9800
```

Against a structural budget of order 0.4 to 0.5 (typical of the worked cases in this manual), a single resource at 99% utilisation can exceed the entire budget by itself. **A named individual who is the sole route to some necessary approval, running at near-full load, is a sustainability-ceiling violation wearing the costume of a productivity success.**

The management rule: **target utilisation of critical shared resources at or below 0.80, and treat any resource above 0.90 as an active fault requiring either capacity or demand reduction.**

This is Ashby's Law of Requisite Variety in a different costume: only variety absorbs variety, and a fully loaded server has no variety left to absorb anything.

5.9 Estimating kappa

`kappa` is the per-active-link cost per period. It is estimable and the procedure is short.

Step 1 — define an active link. Two people constitute an active link in a period if maintaining mutual awareness of each other's work required time in that period. Not "could have talked" — did have to.

Step 2 — count them. For small groups, ask. For larger ones, sample: pick ten people at random, ask each how many colleagues they had to stay synchronised with this period, and estimate the total as `n * mean / 2`.

Step 3 — price them. Estimate hours per period spent on synchronisation, divided by the link count, converted into the same normalised order-units as `Q`. **The unit conversion is the delicate step and it is where estimates go wrong: `kappa` and `eta` and `Q` must all be in one consistent system or the crossover tables are meaningless.** Chapter 24, §24.3 gives the normalisation procedure.

Step 4 — sanity-check the exponent. If you can observe `C` at several group sizes in comparable conditions, regress `log C` on `log n`. A slope near 2 supports the complete-graph postulate; a slope near 1 means your topology is already cellular or hierarchical and you should model `C = kappa_eff * n` instead, with a correspondingly larger viable size.

A useful cross-check that requires no instrumentation: the raw meeting arithmetic.

n	90-min weekly meeting	person-hours/year (46 weeks)	paths
4	6.0 person-hours/wk	276	6
6	9.0	414	15
8	12.0	552	28
12	18.0	828	66
20	30.0	1,380	190
30	45.0	2,070	435

A thirty-person weekly standing meeting consumes 2,070 person-hours a year — more than a full-time post — and it does so to service 435 pairwise relationships of which any given attendee is party to 29. **The all-hands is a device for paying complete-graph cost while delivering star-topology information flow.** It is the worst of both wiring plans and it is the default in most organisations.

5.10 The share of budget consumed by C

The crossover table gives the point where `C` equals the budget. The approach to that point is worth seeing, because it is where the decisions actually get made. With `eta = 0.22`, `kappa = 0.006`, `N = n`:

n	eta∗ln n	C	C as multiple of budget
3	0.2417	0.0180	0.07
5	0.3541	0.0600	0.17
6	0.3942	0.0900	0.23
8	0.4575	0.1680	0.37
10	0.5066	0.2700	0.53
12	0.5467	0.3960	0.72
15	0.5958	0.6300	1.06
18	0.6359	0.9180	1.44
20	0.6591	1.1400	1.73
25	0.7082	1.8000	2.54
30	0.7483	2.6100	3.49
40	0.8116	4.6800	5.77
50	0.8606	7.3500	8.54
100	1.0131	29.7000	29.31
150	1.1023	67.0500	60.83

At `n = 150`, coordination costs sixty-one times the entire structural budget. **Dunbar's number is often quoted as a cognitive limit on stable relationships; it is contested, and this manual states it as contested. What is not contested is the arithmetic: whatever the cognitive limit may be, the economic limit under a complete-graph topology arrives around `n = 15` for these parameters, an order of magnitude sooner.** The reason human groups of 150 exist at all is that they never run a complete graph. They run cells, kin structures, and role hierarchies — topologies, in the sense of §5.7.

5.11 When C grows with Q

A refinement that matters at institutional scale, and one that Chapter 6 depends on.

Several `C` lines are not constant; they scale with accumulated stored order. Compliance reporting (C8) grows with the number of rules to comply with. Maintenance of built artifacts grows with the artifact count. Archive management grows with the archive. Model it as

```
C_t = C_0 + gamma ∗ Q_t
```

Substituting into the storage law:

```
Q_{t+1} = (1 − delta)∗Q_t + eta∗S_t − C_0 − gamma∗Q_t
        = (1 − (delta + gamma))∗Q_t + eta∗S_t − C_0
```

A cost that scales with stored order is arithmetically indistinguishable from a higher decay rate. The system behaves exactly as if `delta_eff = delta + gamma`. Everything from Chapter 2 follows: the time constant shortens to `1/(delta+gamma)`, the steady state falls to `(eta∗S − C_0)/(delta+gamma)`, and the uniform bound of Theorem 3.3 tightens.

Two consequences, both practical.

First, a `Q`-scaling cost line makes an organisation *look* forgetful when it is actually overburdened. The symptom is identical: stored order decaying faster than expected. The remedy is not. Raising retention (Chapter 3) will not touch `gamma`; only reducing the per-unit-`Q` overhead will.

Second, gamma is exactly the term that makes shedding Q a double win, because reducing Q reduces both the delta∗Q maintenance draw and the gamma∗Q cost line at once. Chapter 6, §6.9 develops this into the fourth lever.

5.12 The sustainability ceiling, stated hard

The chapter's conclusion, which is a theorem and should be delivered as one.

```
If   C_t  >  eta ∗ ln N  −  delta ∗ Q_t

then NO arrangement of people prevents decline.
```

Not "decline is likely." Not "decline unless we work harder." The inequality forecloses the entire space of geometric interventions, because ln N is the hard ceiling on F(u) (Paper D, Corollary 2.2) and eta∗ln N is therefore the largest input the interaction channel can possibly deliver in one period. **Every off-site, every reorganisation, every new collaboration model, every diversity initiative, and every leadership change operates strictly inside the space of interventions that Corollary 3.4 has already ruled out.**

There are exactly four responses, and they exhaust the space:

1. **Reduce C_t.** The subject of this chapter.
2. **Raise eta.** The subject of Chapter 4, and usually the cheapest.
3. **Raise N** — logarithmically, therefore weakly, and requiring genuinely distinct categories. Chapter 8.
4. **Reduce Q_t.** Chapter 6, §6.9. The one nobody considers.

Note carefully which response is *not* on the list: raising u. Raising u increases S_t toward its ceiling ln N, which helps — but the ceiling condition above is stated at the ceiling. **If cost exceeds the budget even when S is maximal, then improving the geometry cannot close the gap, and that is precisely what Corollary 3.4 asserts.** Chapters 9 through 11 are essential work, and they are work you do when h > 0, not work you do instead of the arithmetic.

5.13 A note on cutting C badly

Cost reduction has its own failure mode and it deserves stating before the specification.

Every C line is attached to something. Cut C3 (relationship maintenance) and delta rises on trust. Cut C6 (commute and premises) and interaction rate falls, lowering S. Cut C1 (coordination) indiscriminately and you cut the channel through which eta_reach operates, lowering eta. Cut slack and the queueing term explodes (§5.8). **A C cut that raises delta or lowers eta can easily be net negative, and the arithmetic is available to check before the cut rather than after.**

The rule: **for every proposed C reduction, state which term it might damage and by how much, then compute the net effect on h.** A cut of 0.06 in C that costs 0.02 in eta at ln N = 2.08 is a net loss of 0.02∗2.08 − 0.06 = −0.0016. That calculation takes a minute and it is almost never performed.

The C lines that are usually safe to cut are the ones in category C2 — friction — because friction by construction produces nothing. **Attack friction first, coordination second, everything else only with a computed net effect.**

The Specification

C5.1 The system shall maintain a written cost tally decomposed into the nine lines of §5.1. An aggregate `C` figure shall not be accepted as a survey deliverable.

C5.2 The system shall estimate `kappa` by the four-step procedure of §5.9 at least once per twelve periods, and shall record the unit-normalisation used.

C5.3 The system shall compute and record its crossover size $n*$ and its net-optimal size `n_opt = sqrt(eta/kappa)` before any decision that changes headcount.

C5.4 No unit shall be grown past `n_opt` without a recorded justification stating what is being bought with the net structural loss.

C5.5 No unit shall be grown past $n*$ under any circumstances without first changing the topology (§5.7), because past $n*$ the sustainability ceiling is violated at `Q = 0`.

C5.6 The system shall state its communication topology explicitly: which pairs are obliged to maintain a live channel, and which are not. An unstated topology defaults to the complete graph and shall be costed as one.

C5.7 Complete-graph obligation shall be confined to cells of size at most `n_opt`. Inter-cell links shall be explicit, named, and budgeted at an elevated `kappa`.

C5.8 Cell membership shall be permitted to overlap. A partition into disjoint cells with a reporting hierarchy is a tree and shall be flagged as a u-destroying topology (Chapter 8).

C5.9 No topology shall route all inter-cell traffic through a single node. Hub topologies shall be recorded as accepting a known u penalty and a known targeted-failure fragility.

C5.10 Utilisation of every identified critical shared resource shall be measured. Any resource above 0.90 shall be raised as a fault.

C5.11 Target utilisation for critical shared resources shall be at or below 0.80. Slack below that level shall not be reported as waste in any budget document.

C5.12 Every standing recurring meeting shall have its annual person-hour cost computed and displayed on its own agenda.

C5.13 Any standing meeting with more than `n_opt` attendees shall state, in writing, which information flow requires all of them, or be split.

C5.14 The system shall compute the sustainability ceiling `C_max = eta*ln N − delta*Q_t` each period and compare it to measured `C_t`. This comparison shall appear in any report proposing new activity.

C5.15 Where `C_t > eta*ln N − delta*Q_t`, the system shall not authorise any intervention whose mechanism is interaction geometry, and shall record that Corollary 3.4 forecloses it.

C5.16 The system shall test whether any `C` line scales with `Q` (§5.11). Where it does, `gamma` shall be estimated and `delta_eff = delta + gamma` used in all subsequent computation.

C5.17 Every proposed `C` reduction shall be accompanied by a statement of which other parameter it may damage and a computed net effect on `h`.

C5.18 Friction (line C2) shall be attacked before coordination (line C1), and coordination before any other line.

C5.19 Conflict (C4) and disagreement (u) shall be tallied separately and never conflated. A programme that reduces `C4` shall report its measured effect on `u`.

C5.20 No cost reduction shall be implemented that removes a participant's ability to decline an interaction (CANON §5.4). Mandatory attendance is a `C` increase disguised as a `C` control.

The Failure Modes

FM-5.1 . The Quadratic Committee

Damages: `C`. **Description.** A group grown past `n_opt`, and often past `n*`, under complete-graph obligation. Every member is expected to track every other member. **Signature.** Meetings that lengthen without producing more; scheduling as the dominant constraint; the phrase "let's get everyone in a room"; decisions deferred for want of a quorum. **Detection.** Count `n`; compute `sqrt(eta/kappa)`; compare. Compute `C/(eta*ln N)` from §5.10. **Fix.** Cellular topology (C5.7). Split into cells at `n_opt` with explicit liaison links.

FM-5.2 . The Utilisation Cliff

Damages: `C`, via the queueing term. **Description.** A critical shared resource is driven toward full utilisation in pursuit of efficiency. Wait time explodes hyperbolically. **Signature.** Lead times lengthening far out of proportion to load; one named person as the bottleneck for many flows; work-in-progress accumulating; everything "waiting on approval." **Detection.** Measure `rho` on the resource. Anything above 0.90 is the diagnosis. **Fix.** C5.10, C5.11. Add capacity, remove the gate, or shed demand. Restore slack and defend it in the budget as a cost-control device.

FM-5.3 . The All-Hands Trap

Damages: `C`, and `u` indirectly. **Description.** Complete-graph cost paid for star-topology information flow. Everyone attends; one person speaks; nobody's pairwise coordination is served. **Signature.** Large standing meetings with low participation dispersion. Attendees doing other work during them. **Detection.** C5.12: compute the annual person-hour cost. Measure speaking-time distribution (Chapter 14's protocol). **Fix.** C5.13. Split, or convert to a written broadcast, which is star topology at a fraction of the cost.

FM-5.4 . The Uncosted Topology

Damages: every estimate downstream. **Description.** The organisation has never stated who is obliged to coordinate with whom, so everyone assumes everyone, and the cost tally understates `C` badly. **Signature.** Nobody can draw the communication graph. New members ask who they need to keep in the loop and get "everyone, really." **Detection.** Attempt to draw it. The failure to draw it is the finding. **Fix.** C5.6. Write it down. This is a Chapter 4 capture action applied to Chapter 5's problem.

FM-5.5 . The Hub That Ate the Variance

Damages: u, while appearing to fix C. **Description.** Coordination cost reduced by routing everything through a central node. C falls, and so does u, because every message is filtered through one viewpoint. **Signature.** Cheap coordination and unanimous conclusions. Every proposal traceable to one office. Cross-team disagreement never surfaces. **Detection.** Measure u before and after any topology centralisation. Run Chapter 11's disagreement audit. **Fix.** C5.9. Add lateral links that bypass the hub. Accept a higher C in exchange for the u it buys, and compute the trade.

FM-5.6 . The Compliance Ratchet

Damages: C, via gamma; manifests as elevated apparent delta. **Description.** Cost that scales with stored order — reporting, maintenance, archive management — grows unnoticed as Q grows, until delta_eff = delta + gamma dominates. **Signature.** Stored order decaying faster than the measured delta predicts. Rising administrative load with flat output. Older institutions performing worse than younger ones with identical staff. **Detection.** C5.16. Regress C on Q across periods and test whether the slope is significantly positive. **Fix.** Reduce per-unit-Q overhead, or shed Q (Chapter 6, §6.9). Note that shedding Q attacks delta∗Q and gamma∗Q simultaneously.

FM-5.7 . The Conflict-Suppression Trade

Damages: u. **Description.** C4 reduced by discouraging disagreement rather than by improving the handling of it. Cost falls; surplus falls further. **Signature.** Harmonious meetings; low measured conflict; low measured disagreement; decisions unanimously endorsed and privately doubted. **Detection.** C5.19. Tally conflict and disagreement separately. The signature is both falling together. **Fix.** Separate the two. Chapter 14 (acoustics) and Chapter 17 (table geometry) reduce conflict without reducing disagreement, because they remove the physical conditions under which disagreement becomes dominance contest.

FM-5.8 . The Net-Negative Cut

Damages: h, while appearing to raise it. **Description.** A C reduction whose collateral damage to eta or delta exceeds the saving. **Signature.** Cost falls; stored order falls faster. Post-reduction productivity below pre-reduction. **Detection.** C5.17. Compute the net effect on h before cutting. **Fix.** Reverse the cut where possible; this is one of the few cases where restoring C raises h.

FM-5.9 . The Missing Denominator

Damages: the sizing decision. **Description.** Growth justified by total output rather than output per member. Total output can rise while net structural gain falls, throughout the entire band between n_opt and n∗. **Signature.** Headcount growth defended by absolute figures. Nobody computes per-member terms. **Detection.** Compute net(n) = eta∗ln n − kappa∗n(n−1)/2 at current and proposed n. **Fix.** C5.3, C5.4. Make the net computation a required field on any headcount request.

Worked Example

Sizing an incident-response organisation

System. A standing incident-response function inside a mid-size operations group. It must diagnose and resolve heterogeneous failures under time pressure. Its stored order `Q` is measured as the count of live, exercised runbooks and validated diagnostic procedures, normalised so one unit is one such procedure verified within the last quarter.

Survey.

```
eta   = 0.22   (measured per Chapter 4: a strong decision log and runbook discipline)
kappa  = 0.006  (measured per §5.9: 4.2 h/period of synchronisation across 47 active links,
                   normalised to order-units)
delta  = 0.07   (procedures go stale; measured by re-validation failure rate)
Q      = 2.5     (2.5 normalised units of live procedure)
delta*Q = 0.175  (the standing maintenance draw)
```

Every added responder is assumed, generously, to bring a genuinely distinct failure-domain category, so `N = n`. Chapter 7's audit would normally reduce this; taking `N = n` biases the answer toward larger teams, which makes the conclusion conservative.

The sizing table, `h(n) = eta*ln n - delta*Q - kappa*n(n-1)/2`:

n	eta*ln n	C = kappa*n(n-1)/2	h
2	0.1525	0.0060	-0.0285
3	0.2417	0.0180	+0.0487
4	0.3050	0.0360	+0.0940
5	0.3541	0.0600	+0.1191
6	**0.3942**	**0.0900**	**+0.1292**
7	0.4281	0.1260	+0.1271
8	0.4575	0.1680	+0.1145
9	0.4834	0.2160	+0.0924
10	0.5066	0.2700	+0.0616
12	0.5467	0.3960	-0.0243
14	0.5806	0.5460	-0.1404
16	0.6100	0.7200	-0.2850
20	0.6591	1.1400	-0.6559
25	0.7082	1.8000	-1.2668
30	0.7483	2.6100	-2.0367
40	0.8116	4.6800	-4.0434
50	0.8606	7.3500	-6.6644

Results.

```
net-optimal size:            n = 6         (h = +0.1292)
closed form sqrt(eta/kappa): 6.06          (agrees)
h > 0 range:                 n in [3, 11]
C exceeds eta*ln n at:       n = 15
```

Read the table carefully, because it contains three separate boundaries that are routinely confused with each other.

At `n = 2`, `h` is negative. A two-person incident-response function cannot sustain 2.5 units of procedure. There is not enough category diversity to generate the surplus that maintenance requires. **The lower bound is a real bound: understaffing is a sustainability violation exactly as overstaffing is.**

At `n = 6`, `h` peaks. This is the design point.

At `n = 11`, `h` reaches zero. Above eleven, the team cannot maintain the procedures it already has, and Corollary 3.4 applies: no rearrangement of the eleven-plus people prevents `Q` from declining.

At `n = 15`, `C` exceeds the entire structural budget. By then h is deeply negative and the team is consuming its runbooks faster than it writes them.

Note the shape between 6 and 11. h stays positive but falls by 52% across those five hires. **Nothing observable breaks in that band. The team feels busier and produces less durable procedure, and the natural interpretation is that they need more people.** That interpretation is exactly backwards and it is the mechanism by which FM-5.1 propagates.

The eighteen-person question. Suppose the operations group insists on eighteen responders for coverage reasons — a genuine constraint, since incident response requires round-the-clock availability and eighteen bodies may be the minimum to staff a rota. Under a complete graph:

```
C(18)      = 0.006 * 18*17/2 = 0.006 * 153 = 0.9180
h(18)      = 0.22*ln 18 - 0.175 - 0.9180
           = 0.6359 - 0.175 - 0.9180 = -0.4571
```

Deeply unsustainable. But n is fixed by the rota constraint, so three of the four responses in §5.12 are unavailable and the fourth — topology — is the only move.

Restructure as three cells of six with liaison links. Complete graph inside each cell of six; three inter-cell liaison links priced at four times a normal link because they carry translation and handover cost:

```
within-cell:   3 cells * (0.006 * 6*5/2)  = 3 * 0.0900        = 0.2700
between-cell:  3 links * 4 * 0.006        = 0.0720
total C(split) = 0.3420        against C(monolith) = 0.9180

h(split, N = 18) = 0.22*ln 18 - 0.175 - 0.3420
                 = 0.6359 - 0.1750 - 0.3420 = +0.1189
```

The same eighteen people, with the same skills, doing the same work, move from `h = -0.4571` to `h = +0.1189` — a swing of 0.576 — purely by changing who is obliged to keep a live channel with whom. No hiring, no firing, no training, no budget. The cost term falls by 63% and the surplus ceiling is untouched, because `N = 18` regardless of how the eighteen are wired.

Compare that swing against the alternatives:

- Raising `eta` from 0.22 to the level that would rescue the monolith requires `eta' = (delta*Q + C)/ln N = (0.175 + 0.918)/2.8904 = 0.378`, a 72% improvement in conversion efficiency, on top of a discipline that is already strong.

- Cutting `delta` to zero — perfect retention, an impossibility — saves only 0.175 and leaves `h = −0.282`, still negative.
- Shedding `Q` entirely, from 2.5 to 0, saves the same 0.175 and leaves `h = −0.282`.

Only the topology move closes the gap, and it closes it with room to spare. This is the practical content of §5.7 and it is why topology is specified (C5.6–C5.9) rather than left implicit.

What the split costs, stated honestly. The cellular arrangement is not free and the survey must record what was given up.

1. **Cross-cell u may fall.** Responders in cell A stop routinely encountering cell C's failure modes. The countermeasure is overlapping membership (C5.8) — each responder belongs primarily to one cell and secondarily to another, which preserves the semilattice and costs roughly one additional link per person: `18 * 0.006 = 0.108`, taking total `C` to 0.450 and `h` to `+0.0109`. Still positive, and the surplus is protected.
2. **Handover risk rises at cell boundaries.** This is a real `C` line, and the elevated liaison `kappa` of 4x is the estimate of it. If handover proves more expensive — say 8x — total `C` becomes `0.2700 + 0.1440 + 0.108 = 0.5220` and `h = −0.0611`, which is negative. **The sensitivity of the answer to the liaison multiplier is high, and the liaison multiplier is the least well estimated number in this survey.** It should be measured before the split is made permanent, not assumed.
3. **The rota constraint was never tested.** Eighteen was given as a coverage requirement. The arithmetic says a six-person cell is the net optimum, and three cells of six is what the rota produces anyway. The convergence is fortunate rather than designed and should not be assumed to recur in other cases.

Falsification criteria, registered before the split (Chapter 25).

1. Predicted: `Q` rises from 2.5 toward `Q_inf = (eta*S − C)/delta`. With u estimated at 0.55 post-split, `S = ln(1 + 17*0.55) = 2.3370`, `eta*S = 0.5141`, so `Q_inf = (0.5141 − 0.4500)/0.07 = 0.916`. **Note that this is *lower* than the current `Q = 2.5`.** The honest prediction is therefore that the split makes `h` positive while `Q` still falls, because `eta*S` at the achievable u does not support 2.5 units of procedure against `delta = 0.07` plus this `C`. The team must additionally either raise u, raise `eta`, or **shed procedures deliberately** — which is Chapter 6's fourth lever, arriving unprompted from a pure cost analysis. This is the single most useful thing the calculation produced, and it would not have been visible without computing `Q_inf` rather than only `h`.
2. The split is judged a failure if measured `C` after two quarters does not fall by at least half, indicating the liaison multiplier was underestimated (point 2 above).
3. The split is judged a failure if measured u falls below 0.45, indicating the overlapping membership provision failed and FM-5.5 has occurred in cellular form.
4. The split is judged a failure if the shuffled-input control fits the post-split Y series as well as the true pairing, indicating the observed change is a trend rather than a consequence.

Standing caution. `h > 0` licenses building; it does not promise that what exists is sustainable. This example makes both points at once: the split buys headroom, and the same arithmetic reveals that the existing stock of procedure was never supportable at any topology with these parameters. **The cost analysis rescued the organisation and condemned its backlog, and it did both with the same four numbers.**

Chapter 6 . Headroom — The Master Gauge

The Law

Paper E, §6.4, quoted exactly:

```
h_t := eta ∗ ln N  −  delta ∗ Q_t  −  C_t

    h_t < 0    unsustainable in this model. No geometry saves it.
    h_t = 0    maintenance only. Hold, do not build.
  h_t > 0   growth available; any requested increment must satisfy r_t <= h_t.
```

Maximum one-step increment (Corollary 4.2):

```
r_max = eta ∗ ln N  −  delta ∗ Q_t  −  C_t   =  h_t
```

No policy based on interaction geometry alone can guarantee more.

Target increment (Proposition 4.1): `Q_{t+1} − Q_t >= r_t` holds if and only if `S_t >= (delta∗Q_t + C_t + r_t)/eta`, and this is feasible only if `delta∗Q_t + C_t + r_t <= eta∗ln N`.

Headroom is not a new quantity. It is the sustainability ceiling of Corollary 3.4 rearranged into a gauge: `C_t <= eta∗ln N − delta∗Q_t` becomes `h_t >= 0`. **Its value is that it puts three parameters that are usually argued about separately onto one dial with one reading and one sign.** Everything in the previous three chapters — conversion, loss, cost — reports here.

6.0.1 What is proven and what is postulated

Proven. The definition of `h_t` and the three-regime reading follow directly from Paper E, Corollary 3.4 and Corollary 4.2, together with Paper D Corollary 2.2 (`0 <= F(u) <= ln N`), which is what makes `eta∗ln N` an actual ceiling rather than a typical value. Proposition 4.1 is proven. The fixed-point identity derived in §6.6 is a two-line consequence of Theorem 3.5 and requires no additional assumption.

Postulated, badged. That `eta`, `delta`, `N`, and `C` are approximately constant over the horizon on which you read the gauge. That the order proxy `Y_t` you measure stands in a stable relation to `Q_t`. That the linear model is the right model at all. All three are modelling assumptions, all three are falsifiable by Chapter 25's procedure, and **none of them is a theorem.** A headroom figure is only as good as the four measurements behind it, and Chapter 24 is where those measurements are made honest.

Honesty clauses, repeated. The triangle inequality fails for geometric triples (Paper A §5); sub-additivity holds on u only. As deployed, the surplus functional is a lower bound on the ensemble's information — we log the floor, not the ceiling. Paper A is pure mathematics and makes no empirical or metaphysical claim. Paper D does not identify surplus with Shannon entropy or mutual information. Paper E proves no thermodynamic negentropy law: `F(u) + N_can(u) <= ln N` (Corollary 2.2), and higher static surplus does not imply higher canonical concentration.

6.0.2 Why the ceiling term is ln N and not F(u)

A point worth making early because it is the source of most misreadings.

Headroom uses `eta*ln N`, not `eta*F(U_t)`. The difference is deliberate. `F(U_t)` is what the system's *current* geometry delivers this period; `ln N` is the most any geometry could deliver, by Paper D Corollary 2.2, attained only at `u = 1` (perfect orthogonality, Paper A Theorem 4.3).

Headroom is therefore a statement about the best case, and a negative headroom is a statement that even the best case fails. That is what gives Corollary 3.4 its force. When `h_t < 0`, you are not being told that the current arrangement is inadequate. You are being told that the entire space of arrangements is inadequate, and no facilitator, retreat, restructure, or hire inside that space can help. This is the strongest negative result the manual contains, and it is strong precisely because it is stated at the ceiling.

The corollary is that **`h_t > 0` is necessary but not sufficient.** Positive headroom says growth is *permitted*. Whether it *occurs* depends on `F(U_t)`, which is Chapters 7 through 11. The gauge licenses; the geometry delivers.

6.0.3 The gauge as three subtractions

Read the definition as an accounting statement and it becomes usable by anyone:

```
  eta * ln N      the structural BUDGET    -- the most this system can convert in a period
- delta * Q_t      the MAINTENANCE DRAW     -- what must be spent to keep what exists
- C_t            the OPERATING DRAW       -- what is spent to be an organisation at all
 ------------------
 = h_t              HEADROOM                -- what is left over to build with
```

Budget minus two draws. Every organisational conversation about capacity is a conversation about these four numbers, usually with none of them measured.

The Engineering

6.1 Reading the gauge

A headroom reading requires four measurements and one derived quantity. Each has a chapter.

Symbol	Meaning	Where measured	Typical difficulty
`eta`	conversion efficiency	Chapter 4, §4.8	moderate; needs baselines
`N`	count of genuinely distinct categories	Chapter 7	easy to guess, hard to audit
`delta`	retention loss per period	Chapter 3	needs a zero-input window
`Q_t`	stored order, via proxy `Y_t`	Chapter 24	hardest; proxy justification
`C_t`	cost this period	Chapter 5, §5.1	tedious, not hard

The hardest is `Q_t`, and it is hardest for a specific reason that Paper E §6.3 names: **the most common measurement error in the entire framework is choosing an order proxy that is**

really a repackaging of the interaction variable. If Y rises whenever there is a lot of interaction, you have measured U twice and h is meaningless. Chapter 24 gives the test in full.

6.2 The three regimes and what to do in each

Regime 1 — h_t < 0: unsustainable **What it means.** Q will decline this period no matter what. Corollary 3.4 forecloses every intervention whose mechanism is interaction geometry.

What to do. Exactly four things are available, and they exhaust the space:

```
1. raise eta        (Chapter 4)   -- usually cheapest, effect = ln N per unit
2. raise N        (Chapter 8)   -- logarithmic, therefore weak, effect = eta/N per unit
3. cut C            (Chapter 5)   -- effect = 1 per unit; direct
4. shed Q           (this chapter, §6.9) -- effect = delta per unit
```

What NOT to do. Do not commission a workshop, a reorganisation, a diversity initiative, an engagement programme, or a leadership change *as the response to negative headroom*. Each of these operates on F(U_t), and F(U_t) <= ln N already, so the ceiling condition has already ruled them out. They may be excellent things to do once h > 0. They cannot rescue h < 0.

What to say. The correct report is: *"Headroom is negative at -0.12. Under this model, no arrangement of our people prevents decline this period. We must reduce cost, raise conversion, or shed stored order. Here is the arithmetic."* That sentence is worth more than any amount of strategy, because it is checkable.

Regime 2 — h_t = 0: maintenance only **What it means.** The budget exactly covers the two draws. Q can be held, not grown.

What to do. Hold. Refuse new commitments. Protect the maintenance draw absolutely, because at h = 0 any increase in C — a single new reporting requirement, a single new standing meeting — puts the system into Regime 1 immediately.

The characteristic error at h = 0 is to interpret stability as slack. A system holding steady at zero headroom looks calm from outside. It has no reserve at all.

Regime 3 — h_t > 0: growth available **What it means.** An increment r_t up to h_t is achievable. By Corollary 4.2, r_max = h_t exactly, and by Proposition 4.1 it requires S_t >= (delta*Q_t + C_t + r_t)/eta.

What to do. Choose r_t <= h_t, then verify the geometry can deliver the required S_t. If the required S_t exceeds ln N, the increment is infeasible regardless of headroom — this is Proposition 4.1's feasibility clause and it catches the case where the arithmetic permits growth that the geometry cannot supply.

Worked feasibility check. eta = 0.25, N = 8, delta = 0.08, Q_t = 3.0, C_t = 0.10:

```
h_t = 0.25*2.0794 - 0.08*3.0 - 0.10 = 0.5199 - 0.2400 - 0.1000 = +0.1799

request r_t = 0.15  (inside headroom)
required S_t = (0.2400 + 0.1000 + 0.1500)/0.25 = 1.9600 nats
ceiling      = ln 8 = 2.0794
required u   = (e^1.96 - 1)/7 = (7.0993 - 1)/7 = 0.8713
```

Feasible, and only just: it requires `u = 0.871`, which is a very high degree of non-parallelism and is not achievable by asking people to disagree more. The headroom said yes; the feasibility clause said "yes, at a geometry you almost certainly do not have." Requesting `r_t = 0.05` instead requires `S_t = 1.56` and `u = 0.537`, which is attainable. **Always run the feasibility clause after the headroom check. Headroom alone will let you promise things the geometry cannot pay for.**

A concavity correction to the feasibility clause, required by Chapters 24 and 26. The check above computes a single required u as though U were the same every period. It never is, and the error is one-sided. What accumulates is `sum of eta F(U_t)`, so the quantity to test is the **mean of `F(U_t)`, not F of the mean `U`** — and because F is strictly concave (Paper A, derived result 3.3), Jensen's inequality runs against you:

```
mean of F(U_t)  <=  F( mean of U_t )        equality only if U is constant
```

Take this section's parameters, `N = 8`, and a system whose sessions alternate between `u = 0.75` and `u = 0.30`. Mean U is 0.525 and `F(mean U) = 1.5422`, but mean F is 1.4820 — a concavity penalty of 0.0602 nats, **3.9% of the surplus lost to variability alone.** Chapter 26 measures 6.2% in its worked office (0.0761 nats against 1.2262).

Report `U_req` for legibility if you wish; never test against it. State the requirement as `mean F(U_t) >= S_req` and compare it to the measured mean of F. Testing on F of the mean certifies increments the geometry cannot deliver, and the penalty it hides is recoverable for free by regularising a variable geometry. Chapter 24 gives the survey instrument; Chapter 26 §26.7 works the correction end to end.

6.3 Why headroom falls as stored order grows

This is the chapter's central mechanism, and it is one line of algebra:

```
h_t = eta*ln N − delta*Q_t − C_t          =>          d h_t / d Q_t  =  −delta
```

Headroom decreases linearly in stored order at rate `delta`. Every unit of durable structure a system acquires permanently reduces, by `delta` per period, the amount it can build next.

This is the structural reason institutions ossify, and it deserves to be stated without any moral vocabulary at all. **Ossification is not complacency, not cowardice, not bureaucratic capture, and not a failure of leadership. It is `d h / d Q = −delta`.** An institution that succeeds accumulates Q; accumulated Q demands maintenance at `delta*Q`; maintenance consumes the budget; the budget was the capacity to build. **Success consumes its own headroom, by construction, in the absence of any error by anyone.**

The table makes the mechanism concrete. `eta = 0.25`, `N = 8` (budget = `0.5199`), `delta = 0.08`, `C = 0.10`:

Q	maintenance draw `delta*Q`	`h = 0.5199 − delta*Q − 0.10`
0	0.0000	+0.4199
1	0.0800	+0.3399
2	0.1600	+0.2599
3	0.2400	+0.1799
4	0.3200	+0.0999

Q	maintenance draw delta∗Q	h = 0.5199 − delta∗Q − 0.10
5	0.4000	+0.0199
6	0.4800	-0.0601
8	0.6400	-0.2201
10	0.8000	-0.3801
12	0.9600	-0.5401
15	1.2000	-0.7801
20	1.6000	-1.1801

At Q = 0 this system has 0.42 of headroom — it can build freely. At Q = 5 it has 0.02: it is a maintenance organisation. At Q = 6 it is in decline. **Nothing changed except that it succeeded five times.**

The young-institution advantage is now arithmetic rather than folklore. A startup outbuilds an incumbent with more money, more people, and better people, because the startup's delta∗Q is near zero and the incumbent's consumes its entire budget. **The incumbent is not slower. It is spending its budget on being what it already is.**

6.4 The critical stock

Set h_t = 0 and solve for Q:

```
Q_crit  =  ( eta∗ln N  −  C )  /  delta
```

Q_crit is the largest stock of order this system can hold without going into decline. It is a hard capacity, computable from four measured numbers, and it should appear on the front page of any institutional survey.

With eta = 0.25, N = 8 (eta∗ln N = 0.5199):

delta C	0.00	0.05	0.10	0.20	0.40
0.01	51.99	46.99	41.99	31.99	11.99
0.02	25.99	23.49	20.99	15.99	5.99
0.05	10.40	9.40	8.40	6.40	2.40
0.10	5.20	4.70	4.20	3.20	1.20
0.20	2.60	2.35	2.10	1.60	0.60
0.30	1.73	1.57	1.40	1.07	0.40

Two readings.

First, Q_crit scales as 1/delta, so retention is the dominant term in institutional capacity. Moving delta from 0.10 to 0.02 — the work of Chapter 3: documentation, durable artifacts, apprenticeship, ritual, physical permanence — takes capacity from 4.20 to 20.99 at C = 0.10. A factor of five, from retention alone. **The historical institutions that hold the most stored order are not the ones with the most input; they are the ones with the lowest delta.** Written law outlasts oral custom for this reason and no other.

Second, C bites hardest when delta is small. At delta = 0.30, raising C from 0 to 0.40 costs 1.33 units of capacity. At delta = 0.01 the same increase costs 40 units. **The better your retention, the more damaging your overhead** — because low delta means the eta∗ln N − C numerator is divided by a small number, amplifying every change in C.

6.5 The maintenance trap

Named as a standing failure mode (FM-6.2 below) because it has a characteristic and misleading signature.

A system near Q_crit exhibits:

- Full activity, high effort, everyone busy.
- Output that is entirely replacement rather than addition.
- Rising frustration among people who joined to build and find themselves maintaining.
- Any new initiative displacing an existing one rather than adding to it.
- Leadership attributing the stall to insufficient ambition, focus, or talent.

The last item is the trap. **Every symptom of the maintenance trap is also a symptom of a motivation problem, and the two require opposite responses.** Exhortation applied to a headroom problem raises C (more meetings, more reporting, more alignment work) and therefore lowers h, worsening exactly the condition it was meant to fix.

The distinguishing test is arithmetic and takes an afternoon: compute delta∗Q_t and compare it to eta∗ln N. **If the maintenance draw exceeds half the budget, the diagnosis is structural and no amount of motivation is relevant.**

6.6 The fixed-point identity — what headroom a mature system has

Here is a result that follows in two lines from Theorem 3.5 and that changes how the whole chapter reads.

At the steady state with constant U = u∗ and C = c∗:

```
Q_inf = ( eta∗F(u∗) − c∗ ) / delta        (Theorem 3.5)

h_inf = eta∗ln N − delta∗Q_inf − c∗
      = eta∗ln N − ( eta∗F(u∗) − c∗ ) − c∗
      = eta∗ln N − eta∗F(u∗)

h_inf = eta ∗ ( ln N − F(u∗) )
```

At the fixed point, headroom equals the conversion efficiency times the geometric shortfall. delta cancels. C cancels. Q cancels. The only surviving quantities are eta and the gap between what the geometry delivers and what the category count would permit.

Verified numerically across four unrelated parameter sets:

eta	N	u*	delta	C	S = F(u*)	Q_inf	h_inf com- puted	eta*(ln N − S)
0.25	8	0.60	0.08	0.10	1.6487	3.9021	+0.107696	+0.107696
0.18	12	0.45	0.06	0.22	1.7834	1.6835	+0.126273	+0.126273
0.20	14	0.50	0.045	0.08	2.0149	7.1773	+0.124831	+0.124831
0.40	20	0.90	0.03	0.30	2.8959	28.6122	+0.039928	+0.039928

Identical to six decimal places, as the algebra requires.

What this means, and it is the most uncomfortable result in the manual. The only headroom a mature system possesses is exactly the surplus it is *failing* to extract from its own difference. Watch it vanish as the geometry improves, at `eta = 0.25`, `N = 8`:

u*	S = F(u*)	h_inf = eta*(ln N − S)
0.20	0.8755	0.300993
0.40	1.3350	0.186110
0.60	1.6487	0.107696
0.80	1.8871	0.048093
0.90	1.9879	0.022892
0.95	2.0347	0.011184
0.99	2.0707	0.002197
1.00	2.0794	0.000000

A system that has arrived at its fixed point and perfectly exploits its own diversity has zero headroom. At `u = 1` — perfect orthogonality, Paper A Theorem 4.3's extremal case — the surplus is at ceiling, `Q_inf` is at its maximum `eta*ln N / delta` (which is exactly Theorem 3.3's uniform bound, arrived at from a different direction), and there is nothing left over.

This is not a paradox and it is not a criticism of diversity. It is the statement that a fully converged system has, by definition, converged: it has spent its capacity on becoming what it is. Banach guarantees the arrival (CANON §1.3). Nothing guarantees you will like it there. **Book One's Node 10 FIXEDPOINT is Paper E Theorem 3.5, and this identity is the cost of being at it.**

The same zero arrives from the perceptual side, and Chapter 47 gets there first. Volume Two's moment equation `M = (D/P)E` has a fixed point at `D = P` — the calibration ratio `r = 1`, doctrine exactly matching intelligence, undistorted contact. Chapter 47 §47.6 puts that into the storage law and finds `S = ln N`, hence `Q_inf = (eta ln N − C)/delta`, hence

```
h = eta*ln N − delta*Q_inf − C = eta*ln N − (eta*ln N − C) − C = 0    exactly
```

Headroom at perfect calibration is exactly zero by cancellation — which is this section's identity read at `F(u*) = ln N`, arriving from a completely different direction. The consequence is a correction to an obvious-looking goal: **`r = 1` is not an unqualified target.** Undistorted contact maximises what you bank and leaves nothing to build with. An organisation that wants headroom at equilibrium must run below its maximum store — shedding `Q` deliberately (§6.9, Chapter 23) or accepting `r != 1`. See Chapter 47 §47.6 and rule G-12, and Chapter 35 on the ratio itself.

The design consequence is precise and it sets up §6.9:

To retain headroom at maturity you must either hold `eta` high, or refuse to converge, or periodically leave the fixed point by reducing `Q`.

6.7 Why the linear model does not ossify by itself

An honest complication that the previous section forces.

If `h_inf = eta*(ln N - F(u*)) >= 0` always — and it does, since `F(u) <= ln N` by Paper D Corollary 2.2 — then a system running the plain storage law with constant parameters **never reaches negative headroom.** It approaches its fixed point and sits there with a small positive `h`. The trajectory shows it plainly, `eta = 0.25`, `N = 8`, `u = 0.60`, `delta = 0.08`, `C = 0.10`, `Q_0 = 0`:

t	Q_t	h_t
0	0.0000	+0.4199
2	0.5994	+0.3719
5	1.3303	+0.3134
10	2.2070	+0.2433
15	2.7849	+0.1971
20	3.1658	+0.1666
30	3.5822	+0.1333
40	3.7631	+0.1188

Headroom falls monotonically from 0.42 toward its limit 0.1077 and stops. It never crosses zero.

Therefore: observed institutional ossification — real organisations that genuinely cannot build anything — requires something outside the constant-parameter linear model. It is not automatic, and the manual should not pretend it is. Three candidate mechanisms, each with a different remedy, and identifying which one you have is a survey deliverable:

Mechanism A — cost that scales with stored order. `C_t = C_0 + gamma*Q_t` (Chapter 5, §5.11). This is arithmetically identical to a higher decay rate: `delta_eff = delta + gamma`. But note that it does *not* by itself produce negative headroom either, because the fixed-point identity survives the substitution:

```
Q_inf = (eta*S - C_0)/(delta + gamma)
h_inf = eta*ln N - delta*Q_inf - C_0 - gamma*Q_inf
      = eta*ln N - (delta + gamma)*Q_inf - C_0
      = eta*ln N - (eta*S - C_0) - C_0  =  eta*(ln N - S)
```

Identical. A `Q`-scaling cost does not cause ossification; it only makes the system converge to a smaller `Q` faster. This is a genuinely surprising result and it eliminates the most popular explanation.

Mechanism B — exogenous accretion. `Q` arriving through a channel that is not `eta*S`: inherited mandates, merged organisations, statutory obligations, legal accretion, standards adopted by external bodies that you must now uphold. Model it as a constant `A` added each period:

```
Q_{t+1} = (1-delta)*Q_t + eta*S - C_t + A
h_inf   = eta*(ln N - S)  -  A
```

This does go negative, and it goes negative when `A > eta*(ln N − F(u*))`. Exogenous stored order accumulates without paying for its own maintenance, and headroom is what it consumes.

Mechanism C — declining `eta` or `u` with age. Cohesion rising over time drives u toward zero (Chapter 18: cohesion is often the enemy of the collective-intelligence factor c); information cascades collapse diversity in a group that started diverse (Chapter 11); capture ceremonies displace real capture (Chapter 4, FM-4.6). Any of these lowers `eta*S` while `Q` stays high.

Of the three, Mechanism B is the one most often mistaken for laziness and most amenable to a direct fix, and the worked example at the end of this chapter is built on it.

Stated as a single identity for the record, because later chapters lean on it. The fixed-point result of §6.6 is

```
h_inf = eta * ( ln N − F(u*) )        delta, C and Q all cancel
```

and its consequence is the load-bearing half: **the plain linear model never ossifies by itself.** Ossification is not the default state of an ageing institution under this law; it has to be imported, either as exogenous `Q` accretion (Mechanism B) or as a cost that grows superlinearly in `Q` — note that Mechanism A's *linear* `gamma*Q` is not enough, as the cancellation above shows. Chapter 23 takes this up at institutional scope and is where the accretion channel is diagnosed and priced; a system that ossifies is telling you, before any other diagnosis, that it is receiving stored order it did not convert and is not maintaining.

6.8 The four levers, priced

Take the marginal derivative of h with respect to each parameter:

```
d h / d eta  =  ln N          raising conversion is worth ln N per unit
d h / d N    =  eta / N       raising categories is worth eta/N per unit ––
and falls with N
d h / d C    =  −1            cutting cost is worth exactly 1 per unit
d h / d Q    =  −delta        shedding order is worth delta per unit
```

Worked at a common base: `eta = 0.25`, `N = 8`, `delta = 0.08`, `Q = 5`, `C = 0.10`. Baseline h `= 0.5199 − 0.4000 − 0.1000 = +0.0199` — barely above zero, a classic maintenance-trap state.

Lever	Move	New h	Gain	What it costs in practice
L1 raise `eta`	0.25 -> 0.35	+0.2278	+0.2079	a decision log and a recorder rota
L2 raise `N`	8 -> 12	+0.1212	+0.1014	four genuinely new disciplines recruited
L2b raise `N`	8 -> 16	+0.1931	+0.1733	doubling the discipline count

Lever	Move	New h	Gain	What it costs in practice
L3 cut `C`	0.10 -> 0.04	+0.0799	+0.0600	a 60% cut in overhead; usually infeasible
L4 shed `Q`	5.0 -> 3.0	+0.1799	+0.1600	retire two units of stored order
L4b shed `Q`	5.0 -> 2.0	+0.2599	+0.2400	retire three units

Lever 4, shedding three units of stored order, outperforms every other move on the table including a 40% improvement in conversion efficiency. And it is the only one on the list that requires no new capability, no recruitment, no budget, and no negotiation with anyone outside the organisation.

The exchange rate between shedding and the other levers is exact and worth tabulating, because it converts an unfamiliar move into familiar units. To buy headroom increment `dh` by shedding, you must shed `dQ = dh/delta`:

`delta` \ `dh`	0.02	0.05	0.10	0.20	0.50
0.02	1.00	2.50	5.00	10.00	25.00
0.05	0.40	1.00	2.00	4.00	10.00
0.08	0.25	0.62	1.25	2.50	6.25
0.10	0.20	0.50	1.00	2.00	5.00
0.20	0.10	0.25	0.50	1.00	2.50
0.30	0.07	0.17	0.33	0.67	1.67

Note the direction of the effect: shedding is cheapest when `delta` is high. A high-maintenance system gets a large headroom return for a small reduction in stock. A low-`delta` system — one that has invested heavily in permanence — must shed a great deal to recover a little. **Durability and adaptability trade off against each other exactly, and the exchange rate is `1/delta`.** This is the honest form of the intuition that solid institutions are hard to change; it is not a moral fact about institutions, it is `dh/dQ = −delta`.

And the equivalence against `N`, at `eta = 0.25`, `delta = 0.08`:

`N` change	headroom gained	equivalent `Q` to shed
8 -> 10	+0.0558	0.697
8 -> 12	+0.1014	1.267
8 -> 16	+0.1733	2.166
8 -> 32	+0.3466	4.332
8 -> 64	+0.5199	6.498

Retiring 2.2 units of obsolete stored order is worth exactly as much as doubling the number of genuinely distinct disciplines in the organisation. One of those is a decision. The other is a decade.

6.9 The fourth lever — deliberately shedding stored order

This is the lever nobody considers, and it deserves the longest treatment in the chapter.

6.9.1 Why it is invisible Four reasons, and it is worth naming them because each has a countermeasure.

One: `Q` is not measured, so it cannot be reduced. Most organisations have no order observable at all. You cannot decide to shed a quantity you do not track. Countermeasure: Chapter 24's survey.

Two: `Q` is confused with achievement. Stored order is the accumulated record of past success, and reducing it reads as undoing that success. But the law is indifferent to how `Q` was acquired. **A procedure that was excellent in 2011 and is now upheld out of habit costs exactly the same `delta*Q` as one that is excellent today.** The maintenance draw does not inspect provenance.

Three: shedding has an identifiable victim and an unidentifiable beneficiary. The person whose procedure is retired is specific and present. The headroom it releases is diffuse and future. This is a political asymmetry, not an arithmetic one, and it is why shedding must be scheduled in advance rather than negotiated in the moment (§6.9.3).

Four: Star & Ruhleder (1996) — infrastructure becomes visible only on breakdown. Working stored order is invisible. It therefore does not appear on any list of things that could be removed, while its maintenance cost is paid silently every period. Countermeasure: an inventory that lists what exists rather than what is broken.

6.9.2 What can be shed, and what cannot Not all `Q` is equally sheddable, and the distinction should be made explicit before any shedding programme.

Freely sheddable — shed first.

- Procedures whose triggering condition no longer occurs.
- Reports nobody reads. Measurable directly: access logs, or simply stop producing one and count the complaints.
- Approval gates whose rejection rate is near zero. A gate that never rejects is pure `C` and pure `delta*Q`, and it is performing no filtering function at all.
- Superseded documentation retained "for reference." This is FM-4.3, the Write-Only Archive, seen from the maintenance side.
- Committees whose founding question was answered.
- Tooling with no active users.
- Metrics that have been captured by Goodhart's Law and no longer measure what they name.

Conditionally sheddable — shed with a stated replacement.

- Redundant capability. Redundancy is `Q` that buys resilience, which is real. Shed only after computing what failure mode you are re-exposing yourself to. Albert, Jeong & Barabási

(2000) is the relevant caution: robustness to random failure and fragility to targeted attack are different properties, and shedding redundancy affects them differently.
- Institutional memory of past failures. Genuinely valuable and genuinely costly. **Shed the procedure; keep the one-line reason.** A rule with its rationale attached can be re-derived cheaply; a rule without one must be maintained forever because nobody dares remove it.
- Standards commitments to external parties. These require negotiation, not decision.

Not sheddable by this manual, under any headroom pressure.

- Anything constituting a person's ability to decline participation. **CANON §5.4: this manual does not recommend an intervention that removes a person's ability to decline.** A consent provision, a right of appeal, a dissent-protection mechanism, or a minority-report requirement is stored order that will show up in the `delta*Q` draw, and it is not available for shedding regardless of what the arithmetic says. The arithmetic does not have the last word on this and the manual states plainly that it does not.
- Safety interlocks whose failure mode is injury.
- Obligations to parties who are not represented in the decision to shed.

This is the one place in the manual where a lever with a computed positive return is declared out of bounds. It is declared out of bounds anyway.

6.9.3 How to shed: schedule it in advance Shedding negotiated in the moment fails, because §6.9.1's third reason — identifiable victim, diffuse beneficiary — makes every individual case losable. **Shedding works when it is scheduled at the moment of creation, before anyone is attached.** The devices:

- **Sunset clauses.** Every new procedure carries an expiry date. Renewal requires an argument; continuation does not require a decision to end it, which reverses the political asymmetry entirely.
- **Reversal conditions** (Chapter 4, C4.6). Every decision records the observation that would require revisiting it. A decision whose reversal condition has been met and which still stands is a shedding candidate identified by the system itself.
- **Term limits on roles and committees.** The same device applied to organisational structure.
- **Review dates with a default of removal.** The default matters enormously; a review whose default is retention is a formality.
- **Planned demolition.** Jacobs (1961) valued aged buildings for the low rents that let marginal uses survive — a `C`-reduction argument. The complementary point is that a city that never demolishes anything has no site available for anything new. The same is true of a procedure manual.
- **Scheduled forgetting.** Chapter 3's observation, restated: **a system with no forgetting cannot learn.** Shedding is the deliberate form of the `delta` that Chapter 3 spent its length trying to reduce, and the two chapters are not in conflict — Chapter 3 argues for controlling `delta` so it does not remove things at random, and this section argues for removing things on purpose. **Control the leak so you can choose what drains.**

6.9.4 The double return when cost scales with stock Where `C_t = C_0 + gamma*Q_t` (Chapter 5, §5.11), shedding attacks two terms at once:

```
d h / d Q  =  −(delta + gamma)
```

so the shed required to buy headroom dh is dQ = dh/(delta + gamma), which is smaller. **In any institution whose compliance, reporting, or maintenance burden scales with what it has accumulated — which is most large institutions — shedding is strictly more efficient than the §6.8 table suggests.** The worked example quantifies this.

6.9.5 What shedding does not do Honesty requires the limits.

Shedding raises h but does not raise Q_inf. It moves the system away from its fixed point, and Banach guarantees that it will return: the system will re-accumulate toward (eta*S − C_0)/(delta + gamma) at time constant 1/(delta+gamma). **Shedding buys a window, not a new equilibrium.** The window is genuinely valuable — it is the period during which h is large enough to build something structurally different — but a shedding programme that is not paired with a use for the released headroom simply regrows the same stock.

The correct sequencing, which Chapter 26 formalises:

```
1. shed Q          -> h rises, window opens
2. spend h on eta  -> the permanent lever (dh/deta = ln N)
3. re-accumulate at the higher eta -> a higher Q_inf and a higher h_inf
```

Shedding is a transient. Raising eta is permanent. Use the transient to buy the permanent.

6.10 Measuring headroom in practice

The minimum viable headroom reading, achievable in about two weeks in most organisations:

1. **Enumerate N.** List genuinely distinct methodological or domain categories, not job titles, not departments. Chapter 7's collision test: two entries are the same category if they would make the same error on the same problem. Expect the audited N to be roughly half the number first proposed.
2. **Count Q crudely.** A count of live, owned, maintained procedures, artifacts, standards, and standing commitments. Crude is acceptable; the sign of h is usually robust to a factor of two in Q, and the sign is what governs the decision.
3. **Estimate delta** from the last period of genuine inactivity — a shutdown, a holiday closure, a funding gap. delta = 1 − (Q_1/Q_0) per period across that window (Chapter 3).
4. **Tally C** by the nine lines of Chapter 5, §5.1.
5. **Estimate eta** from the last six months of attributable durable change divided by computed S (Chapter 4, §4.8).
6. **Compute h, Q_crit, and h_inf = eta*(ln N − S).** Three numbers on one page.

Report the sign before the magnitude. The sign is what determines which of the four levers you are permitted to reach for; the magnitude only determines how much you may request.

And report the sensitivity. State how much Q would have to be wrong for the sign to flip. If h = +0.02 and a 5% error in Q flips it, say so. Chapter 24, §24.7 gives the sensitivity worksheet.

A cheaper route to the *direction*, derived in Chapter 42. The six-step reading above gives the level of h, and the level is what governs the regime. But the commonest question is not "what is h" — it is "which way should we move." For that there is a shortcut, and it is an identity, not an approximation. Holding eta, N, delta and Q fixed across the configurations being compared, h(x) = eta*ln N − delta*Q − C(x) depends on the configuration only through C, so

```
grad h = −grad C
```

The desire line and the master gauge measure the same gradient, in the same direction, with the same zeros. Book One's Node 02 taught an operator to read where their own choices already go; this identity says that reading *is* a headroom measurement. **A cheap habit survey therefore locates the direction of steepest improvement without the full instrument of Chapter 24** — descending the comfort field is ascending the gauge.

Two cautions, both from Chapter 42 §42.12. First, this gives direction only; the sign of h, which decides whether you may build at all, still requires the four measurements. Second, the identity holds only where `eta`, `N`, `delta` and `Q` really are configuration-independent — and a configuration that changes who is in the room changes `N`. **If a measured desire line points consistently away from measured headroom improvement, that is a diagnosis and not a contradiction: one of those four is moving with the configuration.** The commonest case is low `C` with low `N`, the comfortable narrow life, which is Chapter 11's parallel wall in personal form and is visible only if `N` is measured alongside `C`.

The Specification

C6.1 The system shall compute `h_t = eta*ln N − delta*Q_t − C_t` at least once per twelve periods, and shall record the four measurements and their dates.

C6.2 The sign of `h_t` shall be reported before its magnitude, and shall accompany any proposal for new activity.

C6.3 No increment `r_t` shall be requested such that `r_t > h_t` (Corollary 4.2).

C6.4 Every requested increment shall be checked against Proposition 4.1's feasibility clause: `delta*Q_t + C_t + r_t <= eta*ln N`, and the implied required u shall be computed and stated.

C6.5 Where the implied required u exceeds the measured u by more than a stated margin, the request shall be refused as geometrically infeasible even if `r_t <= h_t`.

C6.6 Where `h_t < 0`, no intervention whose mechanism is interaction geometry shall be authorised. The refusal shall cite Corollary 3.4 and show the arithmetic.

C6.7 Where `h_t < 0`, the response shall be selected from exactly four options: raise `eta`, raise `N`, cut `C`, shed `Q`. Any other proposal shall be recorded as out of scope.

C6.8 The system shall compute and publish `Q_crit = (eta*ln N − C)/delta` and shall report `Q_t / Q_crit` as a standing capacity-utilisation figure.

C6.9 Where `Q_t / Q_crit` exceeds 0.75, the system shall open a shedding review.

C6.10 Where `delta*Q_t` exceeds half of `eta*ln N`, the system shall record that it is in the maintenance regime, and shall not attribute the resulting stall to motivation, ambition, focus, or talent.

C6.11 Every unit of stored order shall carry, at creation, a sunset date or an explicit statement that it is permanent, with the reason.

C6.12 Renewal past a sunset date shall require a positive argument. Continuation shall never be the default at review.

C6.13 The system shall maintain a shedding inventory listing every candidate in the three classes of §6.9.2, updated each review.

C6.14 No item shall be shed that constitutes a person's ability to decline participation, regardless of computed headroom benefit (CANON §5.4).

C6.15 No item shall be shed that constitutes a safety interlock whose failure mode is injury.

C6.16 No item shall be shed that constitutes an obligation to a party not represented in the shedding decision, without that party's participation.

C6.17 Where a procedure is shed, its one-line rationale shall be retained. Shedding the rule while keeping the reason is the specified form.

C6.18 Every shedding programme shall state, in advance, what the released headroom will be spent on. A shedding programme with no stated use shall be deferred (§6.9.5).

C6.19 The system shall test whether any cost line scales with `Q` and, where it does, shall use `d h / d Q = −(delta + gamma)` in all shedding calculations.

C6.20 The system shall compute `h_inf = eta∗(ln N − F(u∗))` and report it as the headroom it will retain once converged. A projected `h_inf` near zero shall be flagged.

C6.21 Every headroom report shall state the sensitivity of the sign of `h` to error in `Q`, `delta`, `eta`, and `C`.

C6.22 No headroom figure shall be published without the order-proxy justification required by Paper E §6.3 and Chapter 4, C6.2's counterpart C4.2.

The Failure Modes

FM-6.1 . The Unread Gauge

Damages: every decision downstream. **Description.** `h` is never computed. Capacity is argued about qualitatively and settled by seniority or volume. **Signature.** Commitments made without any capacity calculation. Repeated overcommitment followed by repeated surprise. "We'll find a way." **Detection.** Ask for the current `h`. The absence of an answer is the finding. **Fix.** C6.1. Two weeks of survey (§6.10) produces a first reading.

FM-6.2 . The Maintenance Trap

Damages: `h`, via `delta∗Q`. **Description.** `Q` has grown until `delta∗Q` consumes the budget. All output is replacement. The stall is attributed to motivation. **Signature.** High activity, low novelty; every new initiative displacing an existing one; builders leaving; exhortation campaigns. **Detection.** C6.10. Compare `delta∗Q_t` to `eta∗ln N`. **Fix.** Shed `Q` (§6.9) or raise `eta` (Chapter 4). Exhortation raises `C` and makes it worse.

FM-6.3 . The Unshed Q

Damages: h, permanently. **Description.** Stored order accumulates monotonically because nothing has a removal path. Every procedure ever adopted is still nominally in force. **Signature.** A procedure manual nobody has read end to end. Rules whose origin nobody knows. Compliance theatre around requirements whose purpose has lapsed. **Detection.** C6.13. Count items with no sunset date. Count approval gates with a rejection rate near zero. **Fix.** C6.11, C6.12. Sunset clauses on everything new; a shedding review on everything old.

FM-6.4 . The Ceiling Misread

Damages: the credibility of the whole method. **Description.** h computed with `eta*F(U_t)` instead of `eta*ln N`, producing a smaller and more pessimistic figure, or conversely `h > 0` read as a promise of growth rather than a licence. **Signature.** Headroom figures that move week to week with meeting quality. Growth promised on headroom alone and not delivered. **Detection.** Check the formula. Check whether the feasibility clause was run. **Fix.** C6.4. h uses the ceiling; feasibility uses the geometry. Both, in that order.

FM-6.5 . The Rising Gauge on a Falling System

Damages: interpretation. **Description.** h rises while Q falls, because the collapsing maintenance draw releases budget. Read as recovery. **Signature.** Improving headroom, deteriorating output, shrinking stored order. Every reading better than the last, and every observable worse. **Detection.** Plot h and Q on the same axes. Divergence in this direction is the diagnosis. **Fix.** Never read h without Q beside it. C6.8 makes `Q/Q_crit` a standing figure precisely to prevent this. The full pattern is worked in Chapter 4's example, where h improves from -0.106 to -0.018 across twenty-six weeks of continuous decline.

FM-6.6 . Exogenous Accretion

Damages: h, through a channel the model does not see by default. **Description.** Q arrives from outside the `eta*S` channel — mandates, mergers, statute, external standards — and consumes headroom without ever having paid for its own maintenance. `h_inf = eta*(ln N - S) - A`, and it goes negative when A exceeds the geometric shortfall. **Signature.** Growing obligations with flat or falling capability. An institution whose remit expands by acquisition rather than by decision. **Detection.** Tally Q acquired by internal decision versus Q acquired by external imposition. Estimate A directly. **Fix.** Shed the accreted Q where permitted (§6.9.2), refuse further accretion, or negotiate the mandate. Note that raising `eta` also helps, since `h_inf = eta*(ln N - S) - A` is increasing in `eta`.

FM-6.7 . The Converged Institution

Damages: h, at the fixed point. **Description.** A system at its steady state with high u has `h_inf = eta*(ln N - S)` near zero. It is performing excellently and has no capacity for anything new. **Signature.** High measured collective performance, high measured diversity, complete inability to take on a new class of problem. Every proposal answered with "we have no capacity." **Detection.** C6.20. Compute `h_inf` directly. Near-zero with high u and high Q is the signature. **Fix.** Raise

eta (the only term that scales h_inf), or shed Q to leave the fixed point deliberately, or accept the state as a legitimate mature equilibrium and say so.

FM-6.8 . Shedding Without a Use

Damages: nothing, and achieves nothing. **Description.** A shedding programme releases headroom that is never spent, and the system re-accumulates the same stock at the same rate. Banach guarantees the return. **Signature.** Periodic "simplification" initiatives whose effects are undone within one time constant. Rule counts that oscillate around a stable mean. **Detection.** Plot Q over several time constants. Oscillation without trend is the finding. **Fix.** C6.18. State the use before shedding, and spend the window on eta (§6.9.5).

FM-6.9 . Shedding Across the Ethical Line

Damages: the standing of the method. **Description.** A shedding programme with a genuine positive computed return removes consent provisions, appeal rights, dissent protections, or safety interlocks, because those are stored order and stored order costs delta∗Q. **Signature.** A headroom argument used to justify removing someone's ability to decline. **Detection.** Check every shedding candidate against C6.14 through C6.16. **Fix.** Refuse. **CANON §5.4 is not subject to the arithmetic, and the arithmetic does not get a vote on it.** Note that the model itself gives no warning here — the computed return is genuinely positive — which is exactly why the constraint must sit outside the model.

Worked Example

A professional standards body over sixty years

System. A professional standards body: a membership organisation that publishes and maintains technical standards for a regulated trade. Annual periods. Q is measured as the count of standards actively maintained — published, in force, and re-validated within the review cycle — normalised so one unit is one such standard.

Survey.

```
eta   = 0.20    conversion of committee work into published, adopted standard
N    = 14     genuinely distinct technical sub-domains (audited per Chapter 7;
                  the initial claim was 31, reduced by the collision test)
u    = 0.50    measured by pre-discussion disagreement on technical questions
delta = 0.045   standards go stale; measured by re-validation failure over a dormant bienniu
C_0  = 0.08    baseline operating cost (governance, publication, membership administration
gamma  =  0.030       cost  that  scales  with  the  standards  count:  re-
validation, cross-reference
                  maintenance, enquiry handling, errata
A    = 0.25    EXOGENOUS ACCRETION: standards adopted by statute and by international
           harmonisation agreements, which the body must maintain and did not choose
```

Derived:

```
S            = ln(1 + 13*0.50) = ln(7.5)  = 2.0149 nats/yr
eta*S        = 0.4030
budget       = eta*ln N = 0.20 * 2.6391    = 0.5278
delta_eff    = delta + gamma               = 0.075
Q_inf      = (eta*S - C_0 + A)/delta_eff = (0.4030 - 0.08 + 0.25)/0.075 = 7.6397
Q_crit      = (budget - C_0)/delta_eff   = (0.5278 - 0.08)/0.075      = 5.9708
h_inf       = eta*(ln N - S) - A        = 0.20*0.6242 - 0.25         = -0.1252
```

The two numbers that decide the case are on adjacent lines: `Q_inf = 7.64` and `Q_crit = 5.97`. The body will converge to a stock 28% larger than the largest stock it can sustain. That is the diagnosis, available in year zero, sixty years before the symptoms.

Trajectory. `Q_{t+1} = 0.955*Q_t + 0.4030 - (0.08 + 0.030*Q_t) + 0.25`:

year	Q_t	C_t	delta*Q_t	h_t
0	0.0000	0.0800	0.0000	+0.4478
5	2.4662	0.1540	0.1110	+0.2628
10	4.1363	0.2041	0.1861	+0.1376
15	5.2672	0.2380	0.2370	+0.0528
20	**6.0331**	**0.2610**	**0.2715**	**-0.0047**
25	6.5518	0.2766	0.2948	-0.0436
30	6.9030	0.2871	0.3106	-0.0699
35	7.1408	0.2942	0.3213	-0.0877
40	7.3019	0.2991	0.3286	-0.0998
45	7.4109	0.3023	0.3335	-0.1080
50	7.4848	0.3045	0.3368	-0.1135
55	7.5348	0.3060	0.3391	-0.1173
60	7.5687	0.3071	0.3406	-0.1198

Headroom crosses zero in year 20, at `Q = 6.033`, just past `Q_crit = 5.971`. Note what year 20 looks like from inside the organisation: it is the twentieth consecutive year of growth, the standards catalogue has grown every year, membership is presumably at a record, and the organisation has by every visible measure succeeded. **There is no event. The crossing has no signature. It is a sign change in a quantity nobody is computing.**

By year 40 the organisation is in a recognisable condition: `delta*Q = 0.329` and `C = 0.299`, together 0.628 against a budget of 0.528. The maintenance draw alone (0.329) is 62% of the total budget, well past the C6.10 threshold. The body publishes little that is new, spends its committee time on revision and errata, and is described by its own members as bureaucratic. That description is accurate and it is not an explanation.

Diagnosis. Run FM-6.6's test. Of the `Q` accumulated, the exogenous accretion term `A = 0.25` per year against internal generation `eta*S = 0.403` per year means **38% of this body's stored order arrived by statute and harmonisation rather than by its own decision, and none of it came with maintenance funding.** The fixed-point identity prices it exactly:

`h_inf = eta*(ln N - S) - A = +0.1248 - 0.2500 = -0.1252`

Without `A`, this organisation would sit at `h_inf = +0.1248` indefinitely — converged, maintained,

modest, and permanently solvent. **The entire deficit is the accretion term.** This is a very different finding from "the organisation became complacent," and it points at a very different remedy.

The four levers, priced at year 60. State: `Q = 7.5687`, `C = 0.3071`, `h = −0.1198`.

L1 — raise `eta`. Required: `eta' = (delta*Q + C)/ln N = (0.3406 + 0.3071)/2.6391 = 0.2454`, from 0.20. **A 23% improvement in conversion efficiency.** Plausible: the standards-drafting process is a capture process and Chapter 4's mechanisms apply directly. This is the best of the first three levers.

L2 — raise `N`. Required: `N' = exp((delta*Q + C)/eta) = exp(0.6477/0.20) = exp(3.2385) = 25.5`, from 14. **The body must acquire eleven more genuinely distinct technical sub-domains — an 82% expansion of its remit.** And note the perversity: expanding remit is precisely the process that generates more `A`. **Lever 2 in this case makes the disease worse while treating the symptom**, which is a structural feature of accretion-driven ossification and not a quirk of these numbers.

L3 — cut `C`. Required: `C' <= budget − delta*Q = 0.5278 − 0.3406 = 0.1872`, from 0.3071 — a 39% cut. But decompose: `C = C_0 + gamma*Q = 0.08 + 0.2271`. **Only 0.08 of the 0.3071 is discretionary; 0.2271 is driven by `Q` itself. Even abolishing the entire baseline operating cost — governance, publication, and administration reduced to zero — leaves `C = 0.2271` and `h = −0.0399`, still negative.** Lever 3 cannot reach the target at any level of austerity. This is the arithmetic behind the familiar spectacle of an institution cutting itself to the bone and remaining stuck.

L4 — shed `Q`. Required: `Q' <= Q_crit = 5.9708`, from 7.5687 — **a 21% reduction, and the deficit is closed.** Because `gamma > 0`, shedding attacks both draws at once at rate `delta + gamma = 0.075` per unit.

Take the shed further, to `0.75 * Q_crit = 4.4781` — retiring 41% of the standards catalogue:

```
new C = 0.08 + 0.030*4.4781            = 0.2143
new h = 0.5278 − 0.045*4.4781 − 0.2143 = +0.1120

headroom recovered = +0.1120 − (−0.1198) = +0.2318
```

A swing of 0.2318, which is 44% of the organisation's entire structural budget, recovered by retiring standards that no longer describe practice. For comparison, achieving the same swing through lever 2 would require `N' = 14 * exp(0.2318/0.20) = 45` — the body would need to triple its technical remit.

Retiring 3.1 units of obsolete standard is worth more than tripling the organisation's technical scope, and it can be decided in one committee meeting.

Execution. The shed is 3.09 units — roughly two fifths of the catalogue — and it is executed against §6.9.2's classification, not by percentage target:

- *Freely sheddable:* standards for equipment classes no longer manufactured; standards superseded by a later standard but never formally withdrawn; guidance notes with no recorded citation in a decade; three committees whose founding question was resolved. **The never-formally-withdrawn category alone is typically the largest and is pure FM-6.3.**
- *Conditionally sheddable, with replacement:* four standards consolidated into one, retaining each original's one-line rationale as an annex (C6.17). This is the "shed the rule, keep the

reason" form and it is what makes the consolidation defensible to the members whose work is being retired.

- *Not shed, and named as not shed:* every statutory obligation (C6.16 — the obligation runs to parties not represented in this decision, so it is renegotiated with the regulator or it stands); every safety-critical interlock standard (C6.15); the members' right of objection to a published standard, which is stored order costing `delta*Q` every year and is **not available regardless of the computed return** (C6.14). **The model gives no warning here. The computed return on removing the objection procedure is genuinely positive. The constraint sits outside the model on purpose.**

Preventing the regrowth. Shedding is a transient (§6.9.5). Banach guarantees the return to `Q_inf = 7.64` at time constant `1/0.075 = 13.3` years, so the released headroom must be spent before it is reabsorbed. The registered plan:

1. Years 61–66: spend `h = +0.112` on lever 1, targeting `eta` 0.20 -> 0.26 through a drafting decision log with mandatory reversal conditions (Chapter 4, C4.4–C4.6). This raises the budget from 0.5278 to 0.6862 and raises `h_inf` from `+0.1248 − A` to `+0.1623 − A`.
2. Year 61 onward: **every new standard carries a sunset date (C6.11) with removal as the review default (C6.12).** This is the structural fix; the shed was the symptomatic one. It converts `Q` from a monotone accumulator into a managed stock and is the only change on this list that prevents recurrence.
3. Years 61–63: open negotiation with the regulator on `A`, seeking either maintenance funding attached to statutory adoptions or a formal withdrawal mechanism. **`A` is the actual cause, and it is the only item on this list that is outside the body's unilateral control.** Every other action treats a symptom, and the survey should say so plainly rather than claim a cure.

Falsification criteria, registered before the intervention (Chapter 25).

1. Predicted: `Q` falls to 4.478 at the shed and, absent the sunset provision, regrows toward 7.64 with time constant 13.3 years — reaching approximately 6.0 by year 74 and re-crossing `Q_crit`. **The sunset provision is falsified if `Q` follows that trajectory anyway.**
2. Predicted: measured `C` falls from 0.3071 to approximately 0.2143 within two review cycles. If `C` does not fall, `gamma` was misestimated and the double-return argument of §6.9.4 does not apply to this body.
3. Predicted: fitted `b` in the minimal model `Y_{t+1} = a*Y_t + b*F(U_t) − C_hat_t + eps_t` rises from 0.20 toward 0.26 following the decision-log intervention. If `b` does not move, the `eta` programme failed and the headroom window will close with nothing bought (FM-6.8).
4. Predicted: fitted `a` remains near `1 − delta_eff = 0.925`. A significant change in `a` indicates that `delta` or `gamma` was not stable across the horizon, which falsifies the constant-parameter assumption badged in §6.0.1 and requires the model to be replaced rather than refitted.
5. The shuffled-input control must degrade the fit. If it does not, the entire fitted relationship between committee interaction and standards output is a trend and the `eta` estimate was never real.

Standing caution. Every number in this example follows from six measured parameters and the theorems of Paper E. **The parameters are the argument; the arithmetic is only bookkeeping.** An organisation that disputes the conclusion should dispute `eta`, `N`, `u`, `delta`, `gamma`, or `A` — and to dispute them it must measure them, which is the entire purpose of putting the gauge on the

wall.

Part II — The Geometry of Difference

Chapter 7 . The Effective Count

Part I built the circuit. We established that stored order obeys a leaky integrator, `Q_{t+1} = (1−δ)Q_t + ηS_t − C_t`, and that everything the engineer can do reduces to moving four parameters: the leak δ, the conversion η, the cost C, and the input S. Part II now opens the last of these. S is not a free variable. It is produced by a geometry, and that geometry has a law.

This chapter is the entry point to that law, and it establishes the single most counterintuitive fact in the manual: **before a system counts its members, the law counts their differences, and a member who differs from no one is not counted at all.**

The Law

7.1 The statement

From Paper A, Theorem 2.1, the interaction surplus of a pair of states x, y is

```
F(x,y) = f(u(x,y))

u(x,y) = 1 − (x . y)^2 = || x ^ y ||^2      in [0,1]

f(u)   = ln(1 + (N−1) u)
```

and the object this chapter is about is the exponential of the surplus:

```
N_eff = g(u) = e^{f(u)} = 1 + (N−1) u        the EFFECTIVE COUNT

surplus = ln N_eff
```

Read the second line slowly. Surplus is not the log of a headcount. **Surplus is the log of a headcount after the headcount has been audited for redundancy.** The audit is performed by u, and u is a purely geometric quantity: it is the squared magnitude of the wedge of the two states, or equivalently one minus the squared cosine between them.

Because `(x . y) = cos θ` for unit states, we may write the identity that will do most of the work in this chapter and the next two:

```
u = 1 − cos^2 θ = sin^2 θ
```

u is the squared sine of the angle between two states. It is 0 when they are parallel and 1 when they are orthogonal, and it moves through the intermediate range in the way a squared sine does — slowly at first, then fast, then slowly again.

7.2 What is proven and what is chosen

The functional form `f(u) = ln(1 + (N−1)u)` is **uniquely characterized** by four axioms. Three of them are bookkeeping. One of them is a design decision, and the manual is required to say so every time it leans on it.

Axiom	Statement	Status
S1	Geometric dependence: F depends on the pair only through u.	Structural. Fixes what kind of object F is.
S2	Zero at zero: `f(0) = 0`. Parallel states produce no surplus.	Normalization. Uncontroversial.
S3	Affine effective count: `g(u) = e^{f(u)}` is affine in u.	**[DESIGN AXIOM — CHOSEN, NOT DERIVED]**
S4	Normalization: `f(1) = ln N`.	Normalization. Sets the ceiling.

S3 is the design axiom. It is chosen, not derived, and it is the reason the effective count is a straight line in u rather than some other curve. Everything in this chapter that feels like a discovery about the world is, at bottom, a consequence of having decided that the effective count should be affine. The decision is a good one — it is the simplest possible assumption that makes "effective count" behave like a count, interpolating linearly between one entity and N entities — but it is a decision. When a reader objects that the second member cannot possibly be worth exactly as much as the fortieth in u-units, the honest answer is: S3 says it is, S3 was chosen, and if you have a measured reason to replace S3 with something else, the rest of the theory will have to be rebuilt around it.

Given S1–S4, the following are **derived**, not assumed (Paper A):

Property	Reference	Statement
Strict monotonicity	Thm 3.1	More non-parallelism strictly increases surplus.
Positivity	Thm 3.2	`F > 0` for any `u > 0`.
Strict concavity	Thm 3.3	Diminishing returns in u, unconditionally.
Subadditivity in u	Thm 3.4	`f(u_1 + u_2) <= f(u_1) + f(u_2)` on the parameter.
Sharp Lipschitz constant	Thm 4.2	`|f(u_1) − f(u_2)| <= (N−1)|u_1 − u_2|`, and N-1 is sharp.
Extremality	Thm 4.3	`F = 0` iff parallel; `F = ln N` iff orthogonal.
Range	Paper D, Cor 2.2	`0 <= F(u) <= ln N`.

7.3 The honesty clauses, restated for this chapter

These are binding wherever the law is used and are reproduced here rather than referenced, because this is the chapter in which a reader is most tempted to over-read the result.

1. **Paper A is pure mathematics and makes no empirical or metaphysical claim.** It proves that a particular functional form follows from four axioms. It does not prove that human

groups obey that form. The empirical content of this manual is a modeling proposal, tested by the protocols in Chapters 24 and 25, and falsifiable there.

2. **As deployed, the surplus functional is a lower bound** on the ensemble's information, not the information itself. We log the floor, not the ceiling. An `N_eff` of 2.5 does not mean the group contains exactly 2.5 independent perspectives; it means the geometry we measured guarantees at least that much, and there may be more that our state encoder failed to see.
3. **Surplus is not identified with Shannon entropy or mutual information in general** (Paper D). The ln is not a Shannon ln. Do not import entropy intuitions wholesale.
4. **The triangle inequality FAILS for geometric triples** (Paper A §5). Subadditivity holds on the parameter u only. You may not chain surplus across three states as though it were a distance. This matters directly in §7.9 when we discuss auditing a roster larger than two.
5. **No thermodynamic negentropy law is proved** (Paper E). `F(u) + N_can(u) <= ln N` bounds a sum; it does not license any statement about physics, order, or consciousness.

7.4 The Roster Reduction — a badged postulate

The law is stated for a **pair** of states x, y, with N entering only through the normalization S4 as the number of distinct categories the ambient space can express. Practical work wants to apply it to a roster of N people, or N teams, or N sensors. That step is not in Paper A. It is a modeling assumption and this manual badges it:

> **THE ROSTER REDUCTION [POSTULATE].** For a body of N members, we take N to be the number of nominal categories the body claims to span, and we take u to be a representative non-parallelism across the body's pairs — normally the mean pairwise u over all `N(N−1)/2` pairs, or a lineage-weighted version of the same. `N_eff = 1 + (N−1)u` is then read as the audited count.

This postulate is doing real work and it can be wrong in a specific way: because the triangle inequality fails for geometric triples, a mean pairwise u does **not** determine the joint geometry of the whole body. Two rosters with the same mean pairwise u can have genuinely different structure — one a ring of mutual near-orthogonality, another a tight cluster plus one distant outlier. The Roster Reduction discards that distinction. It is used because it is cheap, computable from instruments you can actually run (§7.6), and because it errs conservatively in the common institutional case, where the failure is a tight cluster and the mean pairwise u is dragged down by exactly the redundancy we are trying to detect. **Where the difference between "one distant outlier" and "genuine spread" matters to a decision, do not use the mean; report the pairwise matrix and read it directly.** Chapter 10 splits u into its two components and gives a sharper instrument.

The mean-pairwise choice is nevertheless less arbitrary than it looks, and Chapter 24 says why. Let `G` be the Gram matrix of the N encoded unit vectors, `G_ij = x_i . x_j`, and let `PR` be the standard Renyi-2 participation ratio of its spectrum, `PR = (tr G)^2 / ||G||_F^2`. Since `||G||_F^2 = N + N(N−1)(1−u)`, one line of algebra gives

```
N_eff = 1 + (N−1)u = N + 1 − N/PR
```

The mean-pairwise u aggregator IS the participation ratio, reparametrised into effective-count form. The identity is exact; Chapter 24 §24.3 verifies it numerically at `K = 3, 5, 8, 12` on random ensembles to better than 1e-12. Its two poles are worth carrying. At `PR = N` the Gram matrix has no off-diagonal mass, `N_eff = N`, and roster equals effective count — attainable only at exact mutual orthogonality. At `PR = 1` every vector is the same, `N_eff = 1`, and the size of

the roster has left the arithmetic entirely.

This does not repair the postulate; the triangle-inequality objection above still stands, because PR is itself a scalar summary of a spectrum. **What it does is name the audit.** The Roster Reduction is not an improvisation but a standard diversity measure an ecologist, a portfolio analyst or a spectroscopist would recognise on sight, and when a sponsor asks why disagreement is averaged pairwise rather than some other way, the answer is this identity and not a plausibility argument. See Chapter 24 §24.3.

The Engineering

7.5 Ten parallel windows are one window

Set u = 0. Then `N_eff = 1 + (N−1).0 = 1`, and `surplus = ln 1 = 0`.

This is the whole chapter in one line, and it is not a limiting case or an edge condition. It is the modal condition of institutional life. A wall of ten identical windows, all facing the same courtyard at the same angle, admits one view. A panel of ten reviewers trained in the same three departments, reading the same literature, rewarded by the same committee, produces one review. The arithmetic does not care that there are ten signatures on it. It prices the geometry.

The corresponding engineering instruction is blunt: **the first question about any body is not how many members it has but how many directions it points, and the second question is how you measured that.**

Table J gives the audited count directly. Read down a column to see what a nominal roster is actually worth at a given non-parallelism.

TABLE 7.1 — EFFECTIVE COUNT N_eff = 1 + (N−1)u

N	u=0.0	u=0.05	u=0.1	u=0.2	u=0.3	u=0.5	u=0.7	u=0.9	u=1.0
2	1.00	1.05	1.10	1.20	1.30	1.50	1.70	1.90	2.00
3	1.00	1.10	1.20	1.40	1.60	2.00	2.40	2.80	3.00
4	1.00	1.15	1.30	1.60	1.90	2.50	3.10	3.70	4.00
5	1.00	1.20	1.40	1.80	2.20	3.00	3.80	4.60	5.00
6	1.00	1.25	1.50	2.00	2.50	3.50	4.50	5.50	6.00
8	1.00	1.35	1.70	2.40	3.10	4.50	5.90	7.30	8.00
10	1.00	1.45	1.90	2.80	3.70	5.50	7.30	9.10	10.00
12	1.00	1.55	2.10	3.20	4.30	6.50	8.70	10.90	12.00
15	1.00	1.70	2.40	3.80	5.20	8.00	10.80	13.60	15.00
20	1.00	1.95	2.90	4.80	6.70	10.50	14.30	18.10	20.00
25	1.00	2.20	3.40	5.80	8.20	13.00	17.80	22.60	25.00
30	1.00	2.45	3.90	6.80	9.70	15.50	21.30	27.10	30.00
40	1.00	2.95	4.90	8.80	12.70	20.50	28.30	36.10	40.00
50	1.00	3.45	5.90	10.80	15.70	25.50	35.30	45.10	50.00
100	1.00	5.95	10.90	20.80	30.70	50.50	70.30	90.10	100.00

A hundred-person department at u = 0.05 has an effective count of 5.95. Not 100. Not 50. Just under six. If that department's hiring, training, and promotion pipeline is a single funnel, u = 0.05 is not a pessimistic estimate; it is a generous one.

7.6 The geometry of u, in degrees

Because $u = \sin^2 \theta$, the relationship between the intuitive quantity (how differently do these two see the problem?) and the parameter is not linear, and engineers consistently misjudge it in the same direction. Small angles are worth almost nothing.

TABLE 7.2 — ANGLE -> u -> AUDITED COUNT (illustrated at N = 12)

theta	cos θ	u = sin^2 θ	N_eff	ln N_eff	% of ln 12
0deg	1.0000	0.0000	1.000	0.0000	0.0%
5deg	0.9962	0.0076	1.084	0.0802	3.2%
10deg	0.9848	0.0302	1.332	0.2864	11.5%
15deg	0.9659	0.0670	1.737	0.5521	22.2%
20deg	0.9397	0.1170	2.287	0.8271	33.3%
25deg	0.9063	0.1786	2.965	1.0868	43.7%
30deg	0.8660	0.2500	3.750	1.3218	53.2%
35deg	0.8192	0.3290	4.619	1.5302	61.6%
40deg	0.7660	0.4132	5.545	1.7129	68.9%
45deg	0.7071	0.5000	6.500	1.8718	75.3%
50deg	0.6428	0.5868	7.455	2.0089	80.8%
55deg	0.5736	0.6710	8.381	2.1260	85.6%
60deg	0.5000	0.7500	9.250	2.2246	89.5%
70deg	0.3420	0.8830	10.713	2.3715	95.4%
80deg	0.1736	0.9698	11.668	2.4569	98.9%
90deg	0.0000	1.0000	12.000	2.4849	100.0%

Two observations an engineer should carry.

First, **a ten-degree difference of view is worth 3% of an available effective count of twelve.** A team of people who "have different backgrounds" but converge on the same framing within one meeting is a ten-degree team. Everyone experiences the difference as real — they argue, they use different words — and it prices at u = 0.03.

Second, the curve saturates hard past 60deg. Going from 60deg to 90deg — from strongly different to completely orthogonal — buys the last 10.5% of the ceiling. **The expensive, socially painful distance between "strongly disagrees" and "does not share a frame of reference" is worth about a tenth of what the cheap distance between "identical" and "moderately different" is worth.** That is a merciful result and Chapter 9 builds on it heavily: you do not need to import a stranger from another civilization. You need to stop hiring the same person eleven times.

Third, and it is the observation the table's 0deg-to-90deg range conceals: the curve comes back down. u = 1 − (x.y)^2 depends only on the *magnitude* of the correlation, so it is symmetric about 90deg and the table must be read as a fold, not a ramp:

```
120deg  ->  u = 0.7500     same as  60deg
135deg  ->  u = 0.5000     same as  45deg
```

```
170deg  ->  u = 0.0302     same as  10deg
180deg  ->  u = 0.0000     same as   0deg
```

At `rho = −1`, perfect opposition, `u = 0` exactly as it does at `rho = +1`. The member who reliably inverts you spans no more of the space than the member who reliably echoes you: given one of them, you can reconstruct the other, and reconstruction is what redundancy means. **The contrarian carries exactly as much information as the sycophant: none.**

This has a direct consequence for a common staffing move. A body assembled for "creative tension" — a designated opposite, a house sceptic whose position is the negation of the proposal — has bought an `N_eff` of 1 on that axis and is paying `C` for two seats. It is worse than an aligned pair, not better, because it carries the same zero surplus at higher cost. Chapter 9 treats the engineering consequence and Chapter 11 catalogues it as a failure mode distinct from the Parallel Wall.

7.7 Five instruments for estimating u

u is a geometric quantity in a state space you do not directly observe. Estimating it requires an operational choice — Paper E §5.1 calls this the *state encoder*, and Chapter 24 makes the full procedure explicit. Here are the five instruments this manual recommends, in increasing order of cost and decreasing order of ambiguity.

Instrument 1 — Disagreement rate The cheapest instrument. Present the body with a set of K binary or categorical judgment items under **independent** conditions (no discussion, no visible responses, no sequencing). Record the fraction of item-pairs on which two given members differ.

For two members i and j, let `d_ij` be the fraction of items on which they disagree. Under a simple encoding in which each member's response vector is mean-centred and normalized, the sample cosine between response vectors is `1 − 2 d_ij` for balanced binary items, giving

```
u_ij  =  1 − (1 − 2 d_ij)^2
```

`d = 0` gives u = 0. `d = 0.5` — chance-level disagreement, i.e. mutual independence — gives u = 1. `d = 0.1` gives `u = 1 − 0.64 = 0.36`. **Ten percent disagreement is already a third of the way to orthogonality, which tells you how brutally little disagreement most institutions actually tolerate.**

Caveats, and they are severe. The instrument is only valid if elicitation was genuinely independent; a show of hands measures cascade, not geometry (Chapter 11). It is sensitive to item difficulty: items everyone gets right and items everyone gets wrong both read as u = 0 and neither tells you anything about the geometry. Use items with a spread of difficulty and discard items with unanimous responses before computing d.

Instrument 2 — Error decorrelation The most informative single instrument, and the one this manual prefers whenever the body produces predictions that are later scored.

Take each member's **error** vector across a common set of scored items: `e_i = (prediction_i − outcome)` over items. Mean-centre and normalize each error vector. Then the inner product between two normalized, mean-centred vectors is exactly the Pearson correlation coefficient ρ, and

```
u_ij  =  1 − rho_ij^2
```

This is an exact identity, not an approximation, given the encoding choice. **The encoding choice — that error vectors are the right states — is a modeling assumption and is badged as such [POSTULATE]**; the algebra that follows from it is not.

TABLE 7.3 — ERROR CORRELATION -> u -> AUDITED COUNT (illustrated at N = 8)

rho	u = 1 − rho^2	N_eff	ln N_eff	% of ln 8
0.99	0.0199	1.139	0.1304	6.3%
0.95	0.0975	1.683	0.5203	25.0%
0.90	0.1900	2.330	0.8459	40.7%
0.80	0.3600	3.520	1.2585	60.5%
0.70	0.5100	4.570	1.5195	73.1%
0.60	0.6400	5.480	1.7011	81.8%
0.50	0.7500	6.250	1.8326	88.1%
0.40	0.8400	6.880	1.9286	92.7%
0.30	0.9100	7.370	1.9974	96.1%
0.20	0.9600	7.720	2.0438	98.3%
0.10	0.9900	7.930	2.0707	99.6%
0.00	1.0000	8.000	2.0794	100.0%

Note the shape. Because u = 1 - ρ², the audited count is punished by **correlation squared**, which means that the region of high correlation — where institutions actually live — is where the penalty is steepest and least visible. Moving eight forecasters from ρ = 0.90 to ρ = 0.80 nearly doubles the effective count, from 2.33 to 3.52. Moving them from ρ = 0.20 to ρ = 0.10 moves it from 7.72 to 7.93. **The returns to decorrelation are concentrated exactly where organizations feel most comfortable.**

Instrument 3 — Prediction diversity, via a variance identity Standard mathematics, offered here without any citation because none is needed: for N estimators each with error variance σ² and uniform pairwise correlation ρ, the variance of their simple average is

```
Var(mean)  =  (sigma^2 / N) ∗ (1 + (N−1) rho)
```

Define the **classical variance-reduction count** `N_ind = N / (1 + (N−1)ρ)` — the number of independent estimators that would give the same variance reduction. It is worth tabulating because it is the quantity most working analysts already have intuitions about.

TABLE 7.4 — CLASSICAL VARIANCE-REDUCTION COUNT N_ind = N / (1 + (N−1) rho)

N	rho=0.0	rho=0.1	rho=0.2	rho=0.4	rho=0.6	rho=0.8	rho=0.95
2	2.00	1.82	1.67	1.43	1.25	1.11	1.03
4	4.00	3.08	2.50	1.82	1.43	1.18	1.04
8	8.00	4.71	3.33	2.11	1.54	1.21	1.05
12	12.00	5.71	3.75	2.22	1.58	1.22	1.05
20	20.00	6.90	4.17	2.33	1.61	1.23	1.05

```
  50  |   50.00 |    8.47 |    4.63 |    2.43 |    1.64 |    1.24 |    1.05
 100  |  100.00 |    9.17 |    4.81 |    2.46 |    1.66 |    1.25 |    1.05
1000  | 1000.00 |    9.91 |    4.98 |    2.50 |    1.67 |    1.25 |    1.05
```

As N grows, `N_ind → 1/ρ`, a hard asymptote. **A thousand analysts at ρ = 0.2 are worth five independent analysts, and the next nine hundred and ninety are worth nothing at all.**

Now a required warning. `N_ind = N/(1+(N−1)ρ)` and `N_eff = 1+(N−1)u` are **different quantities** and must not be equated. They share a structural motif — the `1 + (N−1).` form — because both are affine interpolations between "one thing" and "N things," which is a natural way to build a count, and S3 chose that same form deliberately. **The resemblance is a design echo, not a theorem, and this manual badges it as such [OBSERVATION, NOT DERIVATION].** Use Table 7.4 for intuition and sanity-checking; use Table 7.3 when you are computing surplus.

Instrument 4 — The collision test From Book One, Node 03. Present two members with the same stimulus under blind conditions and record whether their outputs *collide* — that is, land in the same cell of a pre-declared response partition. Over K trials, the collision rate `c_ij` estimates the probability that the two members are functionally the same instrument.

The collision test is coarser than error decorrelation but it has two advantages that make it the field instrument of choice. It requires no outcome data, so it can be run before anything is scored or on questions that will never be resolved. And it is legible to non-technical participants, which matters more than engineers usually admit: a body will accept a redundancy audit it can see being run, and will contest one that arrives as a correlation matrix from an analyst.

Map collision rate to u by treating the response partition as the encoder:

```
u_ij  ~  1 − c_ij      (for a partition with many roughly equiprobable cells)
```

State the approximation honestly. With a coarse partition — say three cells — random collision alone runs near 1/3, so raw collision rate over-reads similarity. Correct for chance:

```
u_ij  ~  (1 − c_ij) / (1 − c_chance)      clipped to [0,1]
```

where `c_chance` is the collision rate expected if both members answered at random from the observed marginal distribution. **A collision audit that does not subtract chance collision is not a measurement; it is a mood.**

Instrument 5 — The lineage census The bluntest and often the most decisive. Enumerate, for every member, the answers to a fixed set of provenance questions: where they were trained, by whom, in what decade, in what language, under what funding model, with what incentive attached to being right. Two members who answer identically on all axes are provisionally scored at high correlation until measured otherwise.

The lineage census is not a substitute for Instruments 1–4 and this manual does not claim that provenance determines view. **It is a prior, not a measurement, and it is badged as such [HEURISTIC].** Its function is triage: it tells you which pairs to measure first, and it makes the redundancy structure visible to people who will never read a correlation matrix. It is also the only one of the five instruments you can run on a hiring pipeline before anyone is hired, which makes it the only one that can prevent the problem rather than diagnose it.

7.8 Concavity — the merciful clause

`f` is strictly concave in u (Paper A, Thm 3.3), and `ln N_eff` is concave in N as well. The practical consequence is the most useful budgeting fact in the manual.

Fix u. Ask what the Nth member adds:

```
    ΔF(N)  =  ln(1 + (N−1)u)  −  ln(1 + (N−2)u)
```

TABLE 7.5 — MARGINAL SURPLUS OF THE Nth MEMBER, AT FIXED u
(values are nats; the total ceiling is ln N)

N	u = 0.1	u = 0.3	u = 0.5	u = 0.8	u = 1.0
2	0.0953	0.2624	0.4055	0.5878	0.6931
3	0.0870	0.2076	0.2877	0.3677	0.4055
4	0.0800	0.1719	0.2231	0.2683	0.2877
5	0.0741	0.1466	0.1823	0.2113	0.2231
6	0.0690	0.1278	0.1542	0.1744	0.1823
7	0.0645	0.1133	0.1335	0.1484	0.1542
8	0.0606	0.1018	0.1178	0.1292	0.1335
9	0.0572	0.0924	0.1054	0.1144	0.1178
10	0.0541	0.0846	0.0953	0.1027	0.1054
11	0.0513	0.0780	0.0870	0.0931	0.0953
12	0.0488	0.0723	0.0800	0.0852	0.0870
13	0.0465	0.0674	0.0741	0.0785	0.0800
14	0.0445	0.0632	0.0690	0.0728	0.0741
15	0.0426	0.0594	0.0645	0.0678	0.0690
16	0.0408	0.0561	0.0606	0.0635	0.0645
17	0.0392	0.0531	0.0572	0.0597	0.0606
18	0.0377	0.0504	0.0541	0.0564	0.0572
19	0.0364	0.0480	0.0513	0.0533	0.0541
20	0.0351	0.0458	0.0488	0.0506	0.0513
21	0.0339	0.0438	0.0465	0.0482	0.0488
22	0.0328	0.0420	0.0445	0.0460	0.0465
23	0.0317	0.0403	0.0426	0.0440	0.0445
24	0.0308	0.0387	0.0408	0.0421	0.0426
25	0.0299	0.0373	0.0392	0.0404	0.0408
26	0.0290	0.0359	0.0377	0.0388	0.0392
27	0.0282	0.0347	0.0364	0.0374	0.0377
28	0.0274	0.0335	0.0351	0.0360	0.0364
29	0.0267	0.0324	0.0339	0.0348	0.0351
30	0.0260	0.0314	0.0328	0.0336	0.0339
31	0.0253	0.0305	0.0317	0.0325	0.0328
32	0.0247	0.0296	0.0308	0.0315	0.0317
33	0.0241	0.0287	0.0299	0.0305	0.0308
34	0.0235	0.0279	0.0290	0.0296	0.0299
35	0.0230	0.0272	0.0282	0.0288	0.0290
36	0.0225	0.0264	0.0274	0.0280	0.0282

37	0.0220	0.0258	0.0267	0.0272	0.0274
38	0.0215	0.0251	0.0260	0.0265	0.0267
39	0.0211	0.0245	0.0253	0.0258	0.0260
40	0.0206	0.0239	0.0247	0.0252	0.0253

Read the u = 1.0 column. **The second genuinely different member is worth 0.6931 nats. The fortieth is worth 0.0253 — a factor of 27 less, for exactly the same salary.** At u = 0.5 the ratio is 0.4055 to 0.0247, a factor of 16. The pattern is universal across the table because it is concavity, and concavity is a theorem here, not a tendency.

Now read across a row instead. At N = 2, moving from u = 0.1 to u = 1.0 multiplies the gain by 7.3. At N = 40 it multiplies it by 1.23. **Non-parallelism matters enormously for the first few members and barely at all for the last few.** Both readings point the same way and Chapter 26 will use them to rank the levers: get two or three genuinely orthogonal members, then stop buying members and start buying η.

The cumulative form is the one to quote at a budget meeting:

TABLE 7.6 — CUMULATIVE SURPLUS f(N,u) = ln(1 + (N−1)u), SELECTED ROWS

N	u=1.00	u=0.50	u=0.30	u=0.15	u=0.05	u=0.00
2	0.6931	0.4055	0.2624	0.1398	0.0488	0.0000
3	1.0986	0.6931	0.4700	0.2624	0.0953	0.0000
4	1.3863	0.9163	0.6419	0.3716	0.1398	0.0000
5	1.6094	1.0986	0.7885	0.4700	0.1823	0.0000
8	2.0794	1.5041	1.1314	0.7178	0.3001	0.0000
12	2.4849	1.8718	1.4586	0.9746	0.4383	0.0000
20	2.9957	2.3514	1.9021	1.3481	0.6678	0.0000
40	3.6889	3.0204	2.5416	1.9242	1.0818	0.0000
100	4.6052	3.9220	3.4243	2.7632	1.7834	0.0000

7.9 Three honest windows set wide

Compare two cells in Table 7.6.

```
N =  3,  u = 1.00   ->   S = 1.0986
N = 12,  u = 0.15   ->   S = 0.9746
```

Three honest windows set wide capture more than a gallery of twelve pointing the same way. Three people, no shared training, no shared incentive, genuinely orthogonal instruments, beat a twelve-person panel at fifteen percent non-parallelism — and they do it while generating 3.2/2 = 3 communication paths instead of 12.11/2 = 66, which by Chapter 5's cost accounting is a 22-fold reduction in the quadratic term.

The comparison generalizes. Any time you are offered a choice between more members and more difference, compute both cells and compare. The result is almost always the same, because the two terms enter the formula asymmetrically: N enters inside a logarithm, u enters linearly inside that same logarithm and multiplies (N-1). Doubling u is not the same shape of operation as doubling N. Here is the comparison at a concrete baseline:

```
baseline: N = 12, u = 0.15          S = 0.9746
```

```
   double N to 24, u unchanged        S = 1.4929   (+0.5183, twelve new hires)
   hold N = 12, double u to 0.30       S = 1.4586   (+0.4841, zero new hires)
    quadruple N to 48, u unchanged        S = 2.0857    (+1.1111, thirty-
six new hires)
   hold N = 12, u to 0.60              S = 2.0281   (+1.0536, zero new hires)
```

Doubling the effective difference within an existing body is worth almost exactly what doubling the body is worth, and quadrupling the difference is worth almost exactly what quadrupling the body is worth — at a fraction of the cost, and with none of the quadratic coordination penalty. This is the central practical result of Part II and Chapter 9 is devoted to executing it.

7.10 The Lipschitz constant is a measurement-error budget

Paper A Theorem 4.2 gives the sharp Lipschitz constant N-1:

```
| f(u_1) − f(u_2) |  <=  (N−1) | u_1 − u_2 |
```

This is usually read as a smoothness result. Read it instead as a **specification for how precisely you must measure u**. The local slope is

```
df/du  =  (N−1) / (1 + (N−1)u)
```

which equals the bound N-1 only at u = 0 and falls off from there.

TABLE 7.7 — SENSITIVITY df/du, AND THE SHARP BOUND

N	u=0.02	u=0.1	u=0.3	u=0.6	u=1.0	bound (N−1)
3	1.923	1.667	1.250	0.909	0.667	2
5	3.704	2.857	1.818	1.176	0.800	4
10	7.627	4.737	2.432	1.406	0.900	9
20	13.768	6.552	2.836	1.532	0.950	19
50	24.747	8.305	3.121	1.612	0.980	49
100	33.221	9.083	3.225	1.639	0.990	99

The engineering reading: at large N and small u, surplus is violently sensitive to u, so a sloppy estimate of u produces a wildly wrong estimate of S. At N = 12, an uncertainty of ±0.05 in u produces these ranges in S:

```
u = 0.05 ± 0.05   ->   S in [0.0000, 0.7419]    width 0.7419
u = 0.10 ± 0.05   ->   S in [0.4383, 0.9746]    width 0.5363
u = 0.20 ± 0.05   ->   S in [0.9746, 1.3218]    width 0.3472
u = 0.40 ± 0.05   ->   S in [1.5790, 1.7834]    width 0.2044
u = 0.70 ± 0.05   ->   S in [2.0980, 2.2246]    width 0.1266
```

Precisely in the low-u regime where institutions actually operate, a plausible measurement error in u spans nearly the entire range of surplus you were trying to estimate. Do not report a single number for S when u is small and poorly measured. Report the interval. Chapter 25's preregistration protocol makes this a requirement.

7.11 Condorcet says the same thing in another language

The Condorcet jury theorem states that if each of N voters is independently correct with probability $p > 0.5$, the probability that a majority is correct rises with N and approaches 1. The independence clause is not decorative. It is the whole theorem.

Here is the theorem doing its work at $p = 0.6$, computed exactly:

TABLE 7.8 — MAJORITY ACCURACY, N INDEPENDENT VOTERS AT p = 0.6

N = 1	0.6000
N = 3	0.6480
N = 5	0.6826
N = 7	0.7102
N = 9	0.7334
N = 11	0.7535
N = 15	0.7869
N = 21	0.8256
N = 31	0.8716
N = 51	0.9265
N = 101	0.9791

Now break independence. If the voters are perfectly correlated — all one instrument — the majority is correct with probability 0.6 regardless of N. In our language, $u = 0$, $N_{eff} = 1$, surplus = 0, and the roster of 101 is a roster of 1. Condorcet's independence requirement and the surplus law's non-parallelism requirement are **the same requirement stated in two vocabularies.** Condorcet states it as a precondition and then reasons about accuracy; the surplus law states it as a continuous parameter and prices the intermediate cases. Neither is more fundamental. Having both is useful, because Condorcet is the version a lawyer or a board member already believes, and the surplus law is the version you can put in a spreadsheet.

This manual does not claim the two are mathematically equivalent, and asserts no theorem connecting them. **The correspondence is interpretive and is badged [INTERPRETIVE CORRESPONDENCE].**

7.12 Where the effective count enters the circuit

Return to Part I. The storage law is

$$Q_{t+1} = (1-\delta)\, Q_t + \eta\, S_t - C_t, \qquad S_t = \ln(1 + (N-1)\, U_t) = \ln N_{eff}(t)$$

and the master gauge is

$$h_t = \eta \ln N - \delta\, Q_t - C_t$$

Two distinct roles for the count, and confusing them is a recurring error:

- **$\ln N$ is the ceiling.** It appears in headroom, in the sustainability ceiling (Cor 3.4), in the uniform bound (Thm 3.3), in the maximum one-step increment (Cor 4.2). It is what the system *could* draw if every interaction were orthogonal. Chapter 8 is entirely about this term.
- **$\ln N_{eff}$ is the draw.** It is what the system actually delivers this period, and it is strictly less than $\ln N$ unless $u = 1$.

The gap between them, `ln N − ln N_eff = ln(N / N_eff)`, is the **redundancy deficit**. At N = 12 and u = 0.15 it is `2.4849 − 0.9746 = 1.5103` nats, meaning the panel is leaving 61% of its theoretical input on the table every single period, forever, until someone changes the geometry. That deficit compounds through the integrator: over the 6.58-period half-life corresponding to δ = 0.10 (Chapter 3, Table L), it is not a one-time loss but a permanently lowered steady state, since `Q_inf = (ηF(u∗) − c∗)/δ` scales directly with F.

Redundancy is not a missed opportunity. It is a permanent reduction in the fixed point.

The Specification

Numbered, checkable requirements. A body that satisfies EC-1 through EC-12 has been audited for effective count. A body that has not been audited should not be described in a plan as having N members; it should be described as having an unmeasured N_eff.

EC-1 — Declare the encoder. Before any u is estimated, write down in one sentence what a member's *state* is taken to be: their scored error vector, their independent response vector on a declared item set, or their cell assignment under a declared partition. No estimate of u is admissible without a declared encoder (Paper E §5.1).

EC-2 — Elicit independently or do not elicit. All responses used to estimate u must be gathered under conditions where no member can observe another's response before committing their own. Written, simultaneous, sealed. Anything else measures cascade, not geometry. This requirement is absolute and is not subject to convenience.

EC-3 — Report N, u, and N_eff together, never N alone. Any document that states the size of a body must state the audited size beside it, in the form `N = 12, u = 0.15 ± 0.05, N_eff = 2.65 [1.55, 4.30]`. A bare headcount in a design document is a specification defect.

EC-4 — Report an interval, not a point, whenever u < 0.3. Table 7.7 shows the sensitivity. Below u = 0.3 the surplus estimate is dominated by measurement error and a point estimate is misleading.

EC-5 — Subtract chance collision. Any collision-based estimate must state `c_chance` and report the chance-corrected figure. Uncorrected collision rates are not admissible.

EC-6 — Discard unanimous items. Items on which every member agrees carry no information about geometry, whether the unanimity is correct or incorrect. Compute u on the discriminating subset and report how many items were discarded. If more than 80% of items were discarded, the item set is not fit for purpose and must be rebuilt.

EC-7 — Estimate u at least once per δ-half-life. u decays (Chapter 9, Chapter 11). An estimate older than one half-life of the body's convergence process is not a measurement; it is a memory. Chapter 3 gives `t_half = ln 2 / (−ln(1−δ))`; use the convergence rate of u, not the decay rate of Q.

EC-8 — Publish the pairwise matrix when a decision turns on structure. The mean pairwise u is a summary and the Roster Reduction discards structure. When a decision depends on whether difference is spread or concentrated in one member, publish the matrix.

EC-9 — Run a lineage census on every hiring pipeline, annually. Provenance axes at minimum: training institution, principal mentor, decade of formation, working language, funding source, prior industry. Report the entropy of each axis across the current roster. This is a prior, not a measurement (§7.7, Instrument 5), and must be labelled as such wherever reported.

EC-10 — Size against the marginal table, not against a headcount target. Before adding the Nth member, look up the marginal gain in Table 7.5 at your measured u and compare it against the member's full loaded cost including the N−1 additional communication paths they create. If the marginal surplus is below 0.05 nats, the correct move is almost certainly to raise u instead.

EC-11 — Compute and publish the redundancy deficit. `ln N − ln N_eff`. This single number is the most useful summary of a body's geometric health and it belongs on the same dashboard as headcount and budget.

EC-12 — Badge the Roster Reduction wherever it is used. Any report applying `N_eff = 1+(N−1)u` to a roster of people rather than to a declared pair of states must carry the badge and one sentence naming the failure mode: mean pairwise u does not determine joint geometry, because the triangle inequality fails for geometric triples.

The Failure Modes

Each entry: NAME, clinical description, damaged parameter, diagnostic signature, detecting measurement, fix. Chapter 27 consolidates the full catalog; these are the ones native to this chapter.

FM-7.1 . THE HEADCOUNT FALLACY

Damages: the estimate of S, and thereby every downstream budget.

Clinical description. The body reports its size and reasons from it. Plans, budgets, and capability claims are all scaled to N. u is never estimated, so `N_eff` is never known, and the gap between the claimed and delivered input is absorbed silently as underperformance attributed to effort or morale.

Diagnostic signature. Any planning document containing a headcount and no non-parallelism figure. Statements of the form "with twelve reviewers we have comprehensive coverage." Retrospectives that conclude "we need more people" after a body has already grown twice without improving.

Detecting measurement. EC-3 compliance audit. Count planning documents; count the fraction containing a u estimate. Below 100% is a defect.

Fix. Institute EC-3 as a document standard. It costs nothing and it is the single highest-leverage administrative change in this manual, because it forces the question into every document where a decision is made.

FM-7.2 . THE PARALLEL ROSTER

Damages: u -> 0, hence `N_eff -> 1`, hence S -> 0.

Clinical description. Every member is a functional copy of every other. Recruitment, socialization, and promotion have converged the body onto a single instrument. Members experience the condition as agreement, competence, and pleasant working relations. It is the most comfortable of all the failure modes and therefore the most durable.

Diagnostic signature. Disagreement rate under blind elicitation below 5%. Error correlation above 0.9. Meetings that end early. Unanimous votes on non-trivial questions. New members who "fit in immediately."

Detecting measurement. Instrument 2 (error decorrelation) run on the last 30 scored decisions, or Instrument 1 on a fresh blind item set of 20 discriminating items.

Fix. Chapter 9 in full. In outline: recruit across a lineage boundary, install rotation, mandate an adversarial role. Note that the fix is a **u** intervention, not an N intervention; adding members from the same pipeline makes the failure worse by raising C while leaving `N_eff` at 1.

This mode is important enough that Chapter 11 is devoted to its detection and its socially disguised variants.

FM-7.3 . THE MISSING DENOMINATOR

Damages: the credibility of the whole measurement chain.

Clinical description. u is reported without its N. Because `N_eff = 1 + (N−1)u`, a u of 0.4 means an effective count of 1.4 at N = 2 and 40.6 at N = 100. A u reported alone is uninterpretable.

Diagnostic signature. A dashboard tile reading "diversity: 0.42."

Detecting measurement. Inspection.

Fix. EC-3. Never report u without N and `N_eff`.

FM-7.4 . SUBSTITUTION OF PROVENANCE FOR GEOMETRY

Damages: the measurement, not the system — but it hides damage to u.

Clinical description. The lineage census (Instrument 5) is treated as a measurement of u rather than as a prior. The body reports high provenance entropy and concludes it has high u, when in fact its members were drawn from different institutions into a single shared framework and their error vectors correlate at 0.92.

Diagnostic signature. Reports that cite composition and never cite disagreement. Assertion that a body is "diverse" with no scored-error figure anywhere in the document.

Detecting measurement. Run Instrument 2 and compare with what Instrument 5 predicted. A large gap between predicted and measured u is itself the finding, and it is a useful one: it tells you that the homogenizing force is downstream of recruitment, in training or incentive.

Fix. Demote provenance to triage, per EC-9's labelling requirement. Then look downstream: shared training, shared incentive, shared review. Chapter 9 §9.7.

FM-7.5 . THE CASCADE-CONTAMINATED ESTIMATE

Damages: the estimate of u, always in the direction of overstating similarity.

Clinical description. u is estimated from responses gathered in sequence or in view of one another. Members anchored on the first response. The measured disagreement rate reflects the cascade, not the geometry. The body then concludes it is more parallel than it is, or — worse — concludes it is appropriately aligned and stops trying.

Diagnostic signature. Disagreement rate strongly dependent on speaking order. First responder's answer disproportionately represented in the modal response. Round-table elicitation of any kind.

Detecting measurement. Run the same item set twice, once sequentially and once with silent-write-first. Compare the two u estimates. A gap above 0.15 indicates cascade contamination of the sequential protocol; use only the blind estimate thereafter.

Fix. EC-2, absolutely. Silent-write-first, sealed simultaneous submission. Chapter 17 gives the furniture that makes this the default and Chapter 22 gives the protocol.

FM-7.6 . THE SATURATED PANEL

Damages: C, without damaging S — which is why it survives.

Clinical description. A body that has passed the point where additional members buy meaningful surplus keeps adding members. From Table 7.5, at u = 0.5 the 30th member buys 0.0328 nats while adding 29 communication paths. The surplus curve is flat; the cost curve is not.

Diagnostic signature. Growth in headcount with flat or declining output. Meetings where a majority of attendees do not speak. Sub-committees forming spontaneously — which, note, is the body correctly self-organizing away from the saturated configuration.

Detecting measurement. Compute marginal surplus per member at measured u (Table 7.5) and compare against loaded cost per member including the quadratic coordination term from Chapter 5. The crossover point is computable and belongs in every org design review.

Fix. Split. Two bodies of six at u = 0.5 deliver `2 x 1.5041 = 3.008` nats of input against `2 x 15 = 30` communication paths; one body of twelve at the same u delivers 1.8718 nats against 66 paths. **Splitting a saturated body is not an admission of failure; it is the arithmetic.** Chapter 8 §8.9 and Chapter 26 give the sequencing.

Worked Example

The twelve-person review panel

A funding body operates a standing panel that reviews proposals in one technical area. The panel has twelve members. It has operated for six years. Its decisions are scored: three years after each funding round, an independent committee rates each funded proposal's outcome on a five-point scale, and the panel's original scores are compared against that outcome. Six years of data are available.

The body's stated problem: "our decisions have become predictable and our best applicants are going elsewhere." Its proposed solution: expand the panel to thirty members to broaden coverage.

We survey before touching anything.

Step 1 — Declare the encoder (EC-1) State = the member's error vector across scored proposals: for each proposal, the member's original score minus the eventual outcome rating. Six years, 214 scored proposals, twelve members. Vectors mean-centred and normalized per member.

Step 2 — Measure u (Instrument 2) Compute all `12 x 11 / 2 = 66` pairwise Pearson correlations of the normalized error vectors, then `u_ij = 1 − ρ_ij²`.

The lineage census (Instrument 5, run first as triage per EC-9) had already flagged the structure: the twelve members come from three training lineages, in a 5 / 4 / 3 split. The measured correlations confirm the census's prior and sharpen it:

```
within-lineage mean rho = 0.97   ->   u_within = 1 - 0.97^2 = 0.0591
cross-lineage  mean rho = 0.55   ->   u_cross  = 1 - 0.55^2 = 0.6975
```

Pair accounting for a 5 / 4 / 3 split:

```
total pairs        = 12 x 11 / 2                    = 66
within-lineage     = C(5,2) + C(4,2) + C(3,2)
                   = 10 + 6 + 3                     = 19
cross-lineage      = 66 - 19                        = 47
```

Mean pairwise u under the Roster Reduction [POSTULATE]:

```
u_bar = (19 x 0.0591 + 47 x 0.6975) / 66
      = (1.1229 + 32.7825) / 66
      = 33.9054 / 66
      = 0.5137
```

Step 3 — Audit the count

```
N     = 12
u     = 0.5137
N_eff = 1 + 11 x 0.5137 = 6.651
S     = ln 6.651 = 1.8948 nats

ceiling  ln 12 = 2.4849
delivery       = 1.8948 / 2.4849 = 76.3%
redundancy deficit = 2.4849 - 1.8948 = 0.5901 nats
```

This panel is in better shape than it feels. It is delivering 76% of its ceiling. **The problem is not that the panel is parallel; the problem is that the panel is parallel in a specific place, and the mean has hidden where.**

Step 4 — Read the structure, not the mean (EC-8) Publish the block structure. Within lineages, `u = 0.0591`: the five members of the largest lineage are, pairwise, at an effective count of `1 + 4 x 0.0591 = 1.236`. Five people delivering `ln 1.236 = 0.212` nats among themselves. **Those five members are, for review purposes, approximately one member and a quarter.**

Read the panel at lineage level — three categories at cross-lineage non-parallelism:

```
N     = 3 lineages
u     = 0.6975
N_eff = 1 + 2 x 0.6975 = 2.395
S     = ln 2.395 = 0.8734      (ceiling ln 3 = 1.0986, delivery 79.5%)
```

Two readings of the same panel, at two granularities, and the honest summary is: **this is a three-instrument panel with twelve people attached to it, and nine of the twelve are duplicate readings of instruments already present.**

Step 5 — Price the proposed solution The panel proposes growing 12 -> 30. What does that buy? Growth by the existing pipeline recruits into the existing lineages, so the block structure — and therefore `u_bar` — does not change materially. At the panel's own `u_bar = 0.5137`:

```
N = 12   N_eff = 6.651   S = 1.8948   pairs = 66
N = 20   N_eff = 10.76   S = 2.3757   pairs = 190
N = 30   N_eff = 15.90   S = 2.7663   pairs = 435
```

Growth to 30 buys `2.7663 − 1.8948 = 0.8715` nats for 18 new members and a 6.6-fold increase in communication paths. Per member: 0.0484 nats. Consult Table 7.5 — that is the marginal-gain regime where EC-10 says stop.

Worse, growth by the existing pipeline is unlikely to hold `u_bar` constant. Recruiting 18 people into three existing lineages shifts the within/cross pair balance toward within. Recompute with a 12 / 10 / 8 split at 30 members:

```
within pairs = C(12,2) + C(10,2) + C(8,2) = 66 + 45 + 28 = 139
cross  pairs = 435 − 139                                  = 296
u_bar = (139 x 0.0591 + 296 x 0.6975) / 435 = (8.215 + 206.46)/435 = 0.4935
N_eff = 1 + 29 x 0.4935 = 15.31
S     = ln 15.31 = 2.7285
```

Still a gain, but the mean u has fallen and the cost has multiplied. **Growth through an unchanged pipeline dilutes the very quantity it is being purchased to increase.**

Step 6 — Price the alternatives *Option A — rebalance to 4 / 4 / 4 by moving one member out of the largest lineage and one into the smallest.* Two personnel actions, no headcount change.

```
within pairs = 3 x C(4,2) = 18,   cross = 48
u_bar = (18 x 0.0591 + 48 x 0.6975)/66 = (1.064 + 33.48)/66 = 0.5234
N_eff = 1 + 11 x 0.5234 = 6.757
S     = ln 6.757 = 1.9106      gain +0.0158 nats
```

Nearly worthless. Rebalancing within a fixed set of lineages moves 19 pairs to 18. **The lineage *count* is the binding constraint, not the lineage *balance*.**

Option B — add a fourth lineage, 3 / 3 / 3 / 3, by replacing three members. No headcount change.

```
within pairs = 4 x C(3,2) = 12,   cross = 54
u_bar = (12 x 0.0591 + 54 x 0.6975)/66 = (0.709 + 37.665)/66 = 0.5814
N_eff = 1 + 11 x 0.5814 = 7.396
S     = ln 7.396 = 2.0009      gain +0.1061 nats
```

Option C — six lineages of two. Replace roughly five members with people from three new lineages.

```
within pairs = 6 x C(2,2) = 6,   cross = 60
u_bar = (6 x 0.0591 + 60 x 0.6975)/66 = (0.355 + 41.85)/66 = 0.6395
N_eff = 1 + 11 x 0.6395 = 8.034
S     = ln 8.034 = 2.0837      gain +0.1889 nats
```

Option D — twelve distinct lineages, the theoretical limit at this cross-lineage correlation.

```
u_bar = 0.6975    (all pairs are cross-lineage)
N_eff = 1 + 11 x 0.6975 = 8.672
S     = ln 8.672 = 2.1602      gain +0.2654 nats
```

Step 7 — Also attack the cross-lineage correlation Every option above holds ρ_cross = 0.55 fixed. But 0.55 is itself a measured quantity and it is high for people from different training lineages. The likely cause is downstream homogenization: the three lineages read the same journals, attend the same annual meeting, and are scored by the same funder against the same declared criteria. That is a shared incentive acting as a common mode.

Suppose an intervention (Chapter 9) reduces ρ_cross from 0.55 to 0.40, leaving lineage structure at the current 5 / 4 / 3:

```
u_cross_new = 1 − 0.40^2 = 0.84
u_bar = (19 x 0.0591 + 47 x 0.84)/66 = (1.123 + 39.48)/66 = 0.6152
N_eff = 1 + 11 x 0.6152 = 7.767
S     = ln 7.767 = 2.0498      gain +0.1550 nats, zero personnel change
```

Reducing the common-mode correlation by 0.15 outperformed adding a fourth lineage and cost no personnel action at all. This is the shape of result Chapter 9 exists to produce, and it is the reason this manual ranks u interventions above N interventions in Chapter 26's lever ordering.

Step 8 — Put it through the circuit The panel's other parameters, measured per Chapters 3–5: η = 0.45 (decisions produce durable precedent and written criteria at a middling rate), δ = 0.09 per annual cycle (membership rotates, criteria drift, institutional memory is thin), C = 0.55 per cycle (meeting logistics, appeals, conflict-of-interest management, and the appeals load which grows with panel size), and current Q = 9.5 on the declared order proxy.

Current headroom:

```
h = η ln N − δQ − C
  = 0.45 x 2.4849 − 0.09 x 9.5 − 0.55
  = 1.1182 − 0.855 − 0.55
  = −0.2868
```

Negative. The panel is in the regime where, by Corollary 3.4, no arrangement of people prevents decline. This changes the entire analysis. Every option in Step 6 was priced against S, but S is capped at η ln N = 1.1182, and the panel's actual delivery of η S = 0.45 x 1.8948 = 0.8527 is already close to that cap. The panel cannot fix a negative headroom by improving its geometry, because the geometry is nearly maxed against a ceiling set by N and η.

Now price growth to 30 correctly:

```
h(N=30) = 0.45 x ln 30 − 0.09 x 9.5 − C(30)
```

ln 30 = 3.4012, so η ln N rises from 1.1182 to 1.5305, a gain of 0.4123. But C rises too. If the coordination term scales even weakly with pairs — and appeals load, scheduling, and conflict-of-interest management all do — then a 6.6-fold increase in pairs will consume the gain several times over. Taking a modest linear-in-pairs coefficient calibrated to the current panel (C = 0.55 at 66 pairs implies roughly 0.0083 per pair at the margin, though only part of C is pair-driven — say half):

```
C(30) ~ 0.275 + 0.0042 x 435 = 0.275 + 1.827 = 2.10
h(30) = 1.5305 − 0.855 − 2.10 = −1.4245
```

Growing the panel makes the headroom five times worse. This is Chapter 8's thesis arriving early and Chapter 5's crossover calculation doing exactly what it was built to do.

Step 9 — The recommendation Given h < 0, Chapter 26's decision procedure forbids building and requires cutting C, raising η, or shedding Q. In priority order:

1. **Cut C.** Reduce panel size from 12 to 7, choosing the 7 to preserve all three lineages (3 / 2 / 2). Pairs fall from 66 to 21, a 3.1-fold reduction. New u_bar: within = 3 + 1 + 1 = 5, cross = 16, u_bar = (5 x 0.0591 + 16 x 0.6975)/21 = (0.296 + 11.16)/21 = 0.5455. N_eff = 1 + 6 x 0.5455 = 4.273, S = ln 4.273 = 1.4523. Ceiling ln 7 = 1.9459, so η ln N = 0.8757. With C = 0.275 + 0.0042 x 21 = 0.363: h = 0.8757 − 0.855 − 0.363 = −0.342. Still negative — the δQ term is the problem.

2. **Shed Q.** The panel carries 9.5 units of stored order, much of it precedent that no longer discriminates: criteria written for a different research landscape, appeals doctrine accreted over six years, a scoring rubric with sixteen dimensions of which four carry all the variance. Sunset the rubric to four dimensions and formally retire pre-2023 precedent. Estimate the shed at 40%: Q: 9.5 –> 5.7, so δQ: 0.855 –> 0.513. h = 0.8757 − 0.513 − 0.363 = −0.0003. Break-even.

3. **Raise η.** The panel's η = 0.45 is low because decisions produce scores but not written reasoning that binds future rounds. Instituting a two-paragraph written rationale per decision, indexed and searchable, is the classic η intervention from Chapter 4. Estimate η: 0.45 –> 0.62. h = 0.62 x 1.9459 − 0.513 − 0.363 = 1.2065 − 0.876 = +0.3305. **Positive.**

4. **Only now, raise u.** With h = +0.3305, there is room to request an increment. Replace one member with an outsider from a fourth lineage. At 7 members in a 2 / 2 / 2 / 1 split: within = 1+1+1 = 3, cross = 18, u_bar = (3 x 0.0591 + 18 x 0.6975)/21 = (0.177 + 12.555)/21 = 0.6063. N_eff = 1 + 6 x 0.6063 = 4.638, S = ln 4.638 = 1.5343. ηS = 0.62 x 1.5343 = 0.9513 against a ceiling of η ln 7 = 1.2065.

Final position:

```
N = 7,  u = 0.6063,  N_eff = 4.638,  S = 1.5343
eta = 0.62, delta = 0.09, Q = 5.7, C = 0.363
h   = 0.62 x 1.9459 − 0.09 x 5.7 − 0.363 = +0.3305
Q_{t+1} = 0.91 x 5.7 + 0.62 x 1.5343 − 0.363 = 5.187 + 0.951 − 0.363 = 5.775
Q_inf   = (0.62 x 1.5343 − 0.363) / 0.09 = (0.9513 − 0.363)/0.09 = 6.537
```

The panel is smaller, it stores less, and its fixed point is 6.54 rather than the 9.5 it currently carries — which looks like a downgrade until you observe that its current 9.5 is on a declining trajectory ($h < 0$) and its new 6.54 is a stable attractor it converges to from either direction, by Banach.

The panel asked to be made bigger. The arithmetic said: make it smaller, write down your reasoning, throw away half your precedent, and then — only then — hire the outsider.

Step 10 — Falsification (forward reference to Chapter 25) State in advance what would show this analysis wrong.

- If, after the intervention, the measured ρ among the seven does not fall relative to the twelve, the Roster Reduction's block accounting was wrong for this body and the u estimates must be rebuilt from the pairwise matrix.
- If Q on the declared proxy does not stabilize within $3\tau \sim 3/0.09 \sim 33$ cycles — impractically long here, so use the 1-cycle prediction `Q_1 = 5.775 ± measurement error` — the storage-law model is not describing this system.
- If the written-rationale intervention does not produce a measurable rise in η, computed as `ΔQ_attributable / S_input` per Chapter 4, then η was not the binding parameter and the recommendation was wrong in its most expensive component.
- The shuffled-input control (Chapter 25) is mandatory: re-run the fit with the S series randomly permuted in time. If the shuffled fit performs comparably, the apparent relationship between panel geometry and stored order is an artifact and this entire analysis should be discarded.

A framework that cannot fail cannot be trusted, and the fourth bullet is the one that would kill this one.

The effective count is the first gate. It tells you how many instruments a body actually has, which is almost never the number on the roster. The next chapter asks the obvious follow-up: if the audited count is low, can we simply raise the nominal count — add domains, add categories, widen the scope — and buy our way to a higher ceiling?

The answer is a logarithm, and it is not encouraging.

Chapter 8 . Raising N — The Logarithmic Ceiling

Chapter 7 established that a body's audited count is `N_eff = 1 + (N−1)u`, and that the audit is usually brutal. The obvious response is to raise N. Add domains. Widen the remit. Bring in the adjacent discipline, then the one after that. If twelve categories are not enough, take thirty.

This chapter prices that move, and the price is a logarithm.

The result is not a caution. It is a structural prohibition, and it is the single most useful negative result in the manual: **you cannot escape a headroom problem by adding domains.** Not because adding domains is bad, but because the return on N is `ln N`, and `ln N` grows so slowly that the quantities you are trying to outrun — δQ and C, which grow linearly and quadratically — will always win in the end. This chapter proves that with real arithmetic and then tells you what to do instead.

The Law

8.1 The ceiling

From Paper D, Corollary 2.2:

```
0  <=  F(u)  <=  ln N
```

with the upper bound attained if and only if the states are orthogonal (Paper A, Thm 4.3). `ln N` is not a soft target or an asymptote. It is a hard ceiling on surplus per interaction, and it is set entirely by N.

The ceiling enters the theory through axiom **S4**, the normalization `f(1) = ln N`. Read S4 correctly: it says that when two states are maximally non-parallel, the surplus they generate is the log of the number of categories the ambient space can express. N is not the number of people in the room. **N is the number of genuinely distinct categories the system is capable of distinguishing.** Chapter 7's Roster Reduction [POSTULATE] lets us read a roster of people as a set of categories for practical work, but the underlying quantity is a dimension count, not a headcount, and this chapter is about the dimension count.

8.2 Where the ceiling appears

`ln N` does not appear once. It appears in every design inequality in Part I, always multiplied by η, and always as the term that caps what the system can do.

Result	Statement	Role of `ln N`
Paper D, Cor 2.2	`0 <= F(u) <= ln N`	The per-interaction ceiling itself.
Paper E, Thm 3.3	`Q_t <= max{Q_0, η ln N / δ}`	Caps the *entire trajectory* of stored order.
Paper E, Thm 3.3 (δ=0)	`Q_t <= Q_0 + t η ln N`	Caps the growth *rate* when there is no leak.
Paper E, Cor 3.4	Nondecreasing step requires `C_t <= η ln N − δ Q_t`	Caps the affordable cost.

Result	Statement	Role of `ln N`
Paper E, Prop 4.1	Feasible only if `δQ_t + C_t + r_t <= η ln N`	Caps the requestable increment.
Paper E, Cor 4.2	`r_max = η ln N − δQ_t − C_t`	Caps the maximum one-step gain.
Paper E, §6.4	`h_t = η ln N − δQ_t − C_t`	The master gauge.

Five of these seven are the same inequality wearing different clothes, and the clothes matter less than the fact that **η ln N is the entire budget from which δQ and C are paid.** Everything a system can spend on maintenance, coordination, and growth comes out of a single term that scales as the logarithm of its category count.

8.3 The honesty clauses, restated for this chapter

1. **Paper A is pure mathematics and makes no empirical or metaphysical claim.** That `ln N` is a ceiling is a theorem about a functional form. That real institutions are bounded by it is a modeling proposal, falsifiable by the protocols of Chapter 25.
2. **S3, the affine effective count, is a design axiom — chosen, not derived.** The logarithmic ceiling is a consequence of having chosen an affine `g(u)` together with the normalization S4. A different choice of S3 would give a different ceiling law. The manual's negative result about adding domains is therefore conditional on S3, and the conditionality is stated here in the open.
3. **As deployed, the surplus functional is a lower bound**, not the ensemble's information. `ln N` bounds the floor we can log, and the honest claim is that we cannot *guarantee* more than `ln N`, not that no more exists.
4. **Surplus is not Shannon entropy or mutual information** (Paper D). The similarity in form between `ln N` and the entropy of a uniform distribution over N states is a formal resemblance arising from both being normalizations over a category count. **It is not an identification and this manual badges it [FORMAL RESEMBLANCE, NOT IDENTIFICATION].** §8.9 uses it as a bridge to Ashby with the badge attached.
5. **The triangle inequality fails for geometric triples** (Paper A §5). When this chapter discusses composing categories, it composes them through u and the subadditivity relation on u only.
6. **No thermodynamic negentropy law is proved** (Paper E). Nothing here licenses a claim about physical order.

The Engineering

8.4 How slowly the logarithm actually grows

Engineers know that logarithms grow slowly. Almost nobody has looked at the table. Here it is, for every integer from 2 to 100.

`TABLE 8.1 — ln N, N = 2 TO 100`

N	ln N	N	ln N	N	ln N	N	ln N	N	ln N
2	0.6931	22	3.0910	42	3.7377	62	4.1271	82	4.4067
3	1.0986	23	3.1355	43	3.7612	63	4.1431	83	4.4188
4	1.3863	24	3.1781	44	3.7842	64	4.1589	84	4.4308
5	1.6094	25	3.2189	45	3.8067	65	4.1744	85	4.4427
6	1.7918	26	3.2581	46	3.8286	66	4.1897	86	4.4543
7	1.9459	27	3.2958	47	3.8501	67	4.2047	87	4.4659
8	2.0794	28	3.3322	48	3.8712	68	4.2195	88	4.4773
9	2.1972	29	3.3673	49	3.8918	69	4.2341	89	4.4886
10	2.3026	30	3.4012	50	3.9120	70	4.2485	90	4.4998
11	2.3979	31	3.4340	51	3.9318	71	4.2627	91	4.5109
12	2.4849	32	3.4657	52	3.9512	72	4.2767	92	4.5218
13	2.5649	33	3.4965	53	3.9703	73	4.2905	93	4.5326
14	2.6391	34	3.5264	54	3.9890	74	4.3041	94	4.5433
15	2.7081	35	3.5553	55	4.0073	75	4.3175	95	4.5539
16	2.7726	36	3.5835	56	4.0254	76	4.3307	96	4.5643
17	2.8332	37	3.6109	57	4.0431	77	4.3438	97	4.5747
18	2.8904	38	3.6376	58	4.0604	78	4.3567	98	4.5850
19	2.9444	39	3.6636	59	4.0775	79	4.3694	99	4.5951
20	2.9957	40	3.6889	60	4.0943	80	4.3820	100	4.6052
21	3.0445	41	3.7136	61	4.1109	81	4.3944		

Sit with the last cell. A hundredfold increase in category count — from one to a hundred — buys 4.6052 nats. From two to a hundred it buys 3.9120. That is the whole return, for the whole journey, from a small team to a large institution.

Now the same table read as a fraction of what a hundred categories deliver:

TABLE 8.2 — SHARE OF ln 100 CAPTURED BY THE FIRST N CATEGORIES

N	ln N	share of ln 100
2	0.6931	15.1%
3	1.0986	23.9%
4	1.3863	30.1%
5	1.6094	34.9%
7	1.9459	42.3%
10	2.3026	50.0%
15	2.7081	58.8%
20	2.9957	65.1%
30	3.4012	73.9%
50	3.9120	84.9%
75	4.3175	93.8%
100	4.6052	100.0%

Ten categories capture exactly half of what a hundred capture. The remaining ninety categories divide the other half between them. The first five capture 35%. The first three capture nearly a quarter.

This is the same concavity that made Chapter 7's marginal table merciful, and it cuts the same way: the early categories are where the value is, and the late ones are rounding.

8.5 The doubling constant

The most useful single fact about the ceiling is scale-invariance of its increment.

```
    ln(2N) − ln(N)  =  ln 2  =  0.693147...     for every N whatsoever
```

TABLE 8.3 — DOUBLING ALWAYS ADDS ln 2 = 0.6931

N	ln N	gain from N/2 → N
2	0.6931	0.6931
4	1.3863	0.6931
8	2.0794	0.6931
16	2.7726	0.6931
32	3.4657	0.6931
64	4.1589	0.6931
128	4.8520	0.6931
256	5.5452	0.6931
512	6.2383	0.6931
1024	6.9315	0.6931

Doubling the number of domains adds exactly 0.6931 nats, whether you go from 2 to 4 or from 512 to 1024. The absolute gain never changes. The relative gain collapses: 0.6931 is 100% of ln 2 but 10% of ln 1024.

The corollary that governs planning:

```
    to add ONE nat to the ceiling, multiply N by  e = 2.71828...
```

TABLE 8.4 — WHAT +1.0 NAT OF CEILING COSTS IN CATEGORIES

from N = 2	(ln N = 0.6931)	requires N = 5.4
from N = 5	(ln N = 1.6094)	requires N = 13.6
from N = 10	(ln N = 2.3026)	requires N = 27.2
from N = 20	(ln N = 2.9957)	requires N = 54.4
from N = 50	(ln N = 3.9120)	requires N = 135.9
from N = 100	(ln N = 4.6052)	requires N = 271.8

And its inverse, the form to use when someone specifies a required ceiling:

TABLE 8.5 — CATEGORIES REQUIRED FOR A TARGET CEILING

target ln N	N required
0.5000	1.65
1.0000	2.72
1.5000	4.48
2.0000	7.39
2.5000	12.18

3.0000	20.09
3.5000	33.12
4.0000	54.60
4.6052	100.00
5.0000	148.41
6.0000	403.43
7.0000	1096.63

A ceiling of 7 nats requires 1,097 genuinely distinct categories. Not 1,097 people — 1,097 categories, each of which must be distinguishable from every other by the system's own encoder, and each of which must be actually populated by someone capable of occupying it.

8.6 The consequence, stated hard

Here is the structural prohibition this chapter exists to establish.

The master gauge is $h_t = \eta \ln N - \delta Q_t - C_t$. Suppose $h < 0$ — the system is in decline, and by Corollary 3.4 **no arrangement of interaction geometry prevents order from decreasing.** The manager's instinct is to grow: add teams, add disciplines, add scope, add categories, and thereby raise $\eta \ln N$ until h turns positive.

Compute what that requires. To close a deficit of size D by raising N alone:

$$\eta \ln N_{new} - \eta \ln N_{old} = D$$
$$\ln(N_{new} / N_{old}) = D / \eta$$
$$N_{new} = N_{old} \cdot e^{D/\eta}$$

The multiplier is $e^{D/\eta}$, and it is exponential in the deficit.

TABLE 8.6 — MULTIPLIER ON N REQUIRED TO CLOSE A DEFICIT D, AT $\eta = 0.40$

deficit D (nats)	N must be multiplied by $e^{D/0.40}$
0.10	1.3 x
0.25	1.9 x
0.50	3.5 x
1.00	12.2 x
2.00	148.4 x

A deficit of two nats — which in Part I's worked systems is an unremarkable size of shortfall — requires a **148-fold** increase in category count. At $\eta = 0.40$, an institute with eight domains and a two-nat deficit needs 1,187 domains.

And here is the part that makes it a prohibition rather than an expense: **δQ and C do not hold still while you multiply N.** C grows at least linearly in N and, through Brooks' $n(n-1)/2$ communication paths, often quadratically. Q grows because a larger system stores more, so δQ grows too. You are trying to outrun linear and quadratic growth with a logarithm. **The logarithm loses, always, and it loses faster the further you push it.**

This is worth writing as a rule, because it will be violated in almost every organization the reader will ever work in.

> **THE DOMAIN-ADDITION PROHIBITION.** A headroom deficit cannot be closed by raising N. The required multiplier is e^{D/η}, exponential in the deficit, while the cost terms it must outrun grow at least linearly and typically quadratically in N. **If h < 0, adding domains makes h worse.** The admissible responses are: raise η, cut C, or shed Q. Chapter 6 named these; this chapter proves the fourth option is not an option.

Two honest qualifications, because the prohibition is strong and strong claims need their limits stated.

First, the prohibition applies to raising N as a **remedy for a deficit**. Raising N for other reasons — because the environment genuinely presents variety the system cannot currently absorb (§8.9), because a new category is cheap and the ceiling is close to binding, because you are building the initial system rather than repairing one — remains legitimate. §8.11 gives the four legitimate cases.

Second, the prohibition assumes η is held constant while N grows. In practice adding a genuinely new domain sometimes raises η, because a new discipline brings new capture mechanisms — a new notation, a new archive, a new form of artifact. That effect is real and this manual does not deny it. But it should be attributed correctly: the gain came from η, not from N, and if you want it you should go get η directly rather than buying a domain and hoping. Chapter 4 gives the direct route and it is far cheaper.

8.7 What raising N actually requires

Suppose you have decided, legitimately, to raise N. The question is what counts.

N is the number of **genuinely distinct categories the system can express**. It is not the number of labels in the org chart. Splitting a category into two relabeled halves does not raise N; it raises the label count while leaving the geometry untouched, and the arithmetic of Chapter 7 prices the result at zero.

Make this precise. Suppose the system has N categories and you split one of them into two halves with pairwise non-parallelism ε. Two things happen:

- The **nominal ceiling** rises from `ln N` to `ln(N+1)`.
- The **delivered surplus** does not, because the new pair contributes u = ε ~ 0.

TABLE 8.7 — CEILING GAIN FROM ADDING ONE CATEGORY

N -> N+1	ceiling gain ln(N+1) − ln N	as % of current ceiling
2 -> 3	0.4055	58.50%
3 -> 4	0.2877	26.19%
4 -> 5	0.2231	16.10%
5 -> 6	0.1823	11.33%
6 -> 7	0.1542	8.60%
8 -> 9	0.1178	5.66%
10 -> 11	0.0953	4.14%
12 -> 13	0.0800	3.22%
16 -> 17	0.0606	2.19%
20 -> 21	0.0488	1.63%

```
 32  ->  33          0.0308                  0.89%
 50  ->  51          0.0198                  0.51%
100  -> 101          0.0100                  0.22%
```

At N = 20, one more category raises the ceiling by 1.63%. If that category is a relabeled half of an existing one, the delivered surplus rises by approximately nothing while C rises by the cost of running a twenty-first category. Category inflation is a pure cost with a cosmetic ceiling gain, and it is the standard institutional response to a demand for breadth.

The relabeling audit Before N is incremented in any plan, run this test on the proposed new category:

1. **State the new category's encoder.** What state does a member of this category occupy that a member of no existing category occupies? Write one sentence. If you cannot, stop.
2. **Estimate u between the new category and its nearest existing neighbour** using Chapter 7's Instruments 1 or 2. Blind elicitation on 20 discriminating items is sufficient.
3. **Require u >= 0.3 against every existing category.** Below 0.3 the new category is a relabel and must not be counted in N. Report the measured u alongside the request.
4. **Compute delivered gain, not ceiling gain.** The plan must state `Δ(delivered) = ln(1 + N.u_bar_new) − ln(1 + (N−1).u_bar_old)`, not `ln(N+1) − ln N`.
5. **Compute the cost delta.** New coordination paths, new meetings, new interfaces. Chapter 5's accounting.
6. **Compare.** If delivered gain is below cost delta, the category is a liability.

Steps 1 and 3 kill most proposals, which is the point. **A category that cannot state what it sees that nobody else sees is a label, and labels do not enter the law.**

The three ways N is faked

- **Splitting.** One category becomes two with the same training, the same literature, the same incentives, and a boundary drawn on the org chart. Diagnostic: measured u between the halves below 0.15. This is the most common.
- **Importing the adjacent.** Recruiting from the discipline next door — which shares 80% of its methods, its journals, and its reviewers with yours. Genuine but small u. Diagnostic: measured u between 0.15 and 0.35, and error correlation above 0.8. Counts partially; count it partially.
- **Nominal representation.** Adding a category to the chart and staffing it with someone whose actual working practice is unchanged from their previous category. Diagnostic: the person's error vector is unchanged from before the reassignment. This is the hardest to detect and requires scored longitudinal data.

8.8 The scale-free reading

The doubling constant has an architectural consequence worth stating, because it is where CANON §1.4's hyperrecursive symmetry becomes concrete.

Because `ln(2N) − ln N = ln 2` at every scale, the ceiling gain from doubling is identical at a desk, in a room, in a team, in a building, in an institution, and in a city. Two drawers instead of one; two disciplines instead of one; two wings instead of one; two boroughs instead of one. All worth 0.6931.

That is a genuine scale-free property of the model form, and it is honest to state it as such: **it is a claim about the FORM of the equation, not about physics.** The parameters δ, η, C differ enormously across those scales; only the shape of the N-dependence is invariant. This is the architectural claim of the manual and it is the reason a single set of design rules can be written for objects that differ in size by six orders of magnitude.

It also explains a common and correct intuition: that adding the second of anything feels transformative and adding the fiftieth feels like paperwork. The absolute gain is the same 0.6931 per doubling either way, but the doubling from 1 to 2 is one hire and the doubling from 50 to 100 is fifty.

8.9 Ashby's Law of Requisite Variety — why N matters at all

Everything above has been a case against raising N. Ashby supplies the case for.

Ashby's Law of Requisite Variety (1956): only variety can absorb variety. A regulator can hold a system's essential variables within limits only if the regulator's variety is at least as great as the variety of the disturbances it must counter. A thermostat with two states cannot regulate a process with ten distinguishable failure modes. A committee with one framework cannot govern a domain that presents four kinds of problem.

Ashby measured variety as **the logarithm of the number of distinguishable states**. That is his own formulation, and it is why the correspondence with `ln N` is close enough to be useful:

```
Ashby:        variety V = log (number of distinguishable states)
Surplus law:  ceiling   = ln N,  N = number of distinguishable categories
```

TABLE 8.8 — VARIETY IN NATS AND BITS

distinguishable states	ln (nats)	log2 (bits)
2	0.6931	1.0
4	1.3863	2.0
8	2.0794	3.0
16	2.7726	4.0
64	4.1589	6.0
256	5.5452	8.0
1024	6.9315	10.0

This manual badges the correspondence [FORMAL CORRESPONDENCE, NOT IDENTIFICATION]. Ashby's variety and the surplus ceiling are both logarithms of a state count, and both function as caps on regulatory capability. They are not proved equal, and Paper D explicitly declines to identify surplus with information-theoretic quantities. Use the correspondence as a bridge between two literatures, not as a theorem.

With the badge attached, the pairing is powerful, because Ashby and the ceiling say opposite and complementary things:

- **Ashby tells you the requirement.** If the environment presents V_env of variety and your regulator has V_reg < V_env, the shortfall shows up as unregulated variation — surprises, crises, failures the system did not have a category for. You cannot regulate what you cannot distinguish, and no amount of effort substitutes for a missing category.

- **The ceiling tells you the price.** Each doubling of your regulatory repertoire costs the whole effort of doubling and returns 0.6931 nats.

Put together: N must be raised when and only when the environment presents variety the system cannot distinguish, and it must be raised by acquiring genuinely new distinctions rather than new labels — because Ashby's requirement is about distinguishability, which is exactly what §8.7's relabeling audit tests.

This gives the four legitimate reasons to raise N, listed in §8.11.

Ashby also explains why the Domain-Addition Prohibition does not contradict him. Ashby's law is about **capability**, not about **budget**. A system that lacks requisite variety will fail at regulation no matter how much headroom it has. A system with requisite variety and negative headroom will fail at persistence no matter how well it regulates. They are different failure modes with different fixes, and confusing them produces the characteristic institutional error: responding to a persistence problem ($h < 0$) with a capability intervention (add domains), which worsens the persistence problem without addressing anything.

Diagnose which one you have before you act. Ashby's failure looks like surprise. The headroom failure looks like erosion.

8.10 Alexander: a city is not a tree

Christopher Alexander's 1965 essay "A City is Not a Tree" is the most directly relevant piece of design theory in this manual's lineage, and Chapter 8 is where it belongs, because it is an argument about **category structure** and therefore about N.

Alexander's distinction. A **tree** is a collection of sets in which any two sets are either disjoint or one contains the other. A **semilattice** is a collection of sets in which two sets may also **overlap** — sharing members without either containing the other. Every tree is a semilattice; the converse is dramatically false.

Alexander's arithmetic, which anyone can verify:

TABLE 8.9 — DESCRIBABLE SUBSETS ON n ELEMENTS

n	tree: at most (n−1) non-trivial subsets	ln	semilattice: 2^n − n − 1	ln	ceilings
4	3	1.0986	11	2.3979	2.18 x
6	5	1.6094	57	4.0431	2.51 x
8	7	1.9459	247	5.5094	2.83 x
10	9	2.1972	1013	6.9207	3.15 x
12	11	2.3979	4083	8.3146	3.47 x
16	15	2.7081	65519	11.0901	4.10 x
20	19	2.9444	1048555	13.8629	4.71 x

On twenty elements a tree can carry at most nineteen further subsets. A semilattice on the same twenty can carry 1,048,555. Alexander stated this in the essay as "more than a million," and the count is exactly `2^20 − 20 − 1`.

Translated into this manual's terms: the tree structure caps N at n-1, while the semilattice structure caps it at 2^n - n - 1, and since the ceiling is ln N, the tree costs you a factor of 4.71 in ceiling at twenty elements and the factor grows with n.

But the deeper point is not about the ceiling. It is about u, and it is the reason the manifest calls hierarchy the organizational form that drives u toward zero.

Why hierarchy drives u to zero In a tree, every element belongs to exactly one chain of nested sets: itself, its parent, its grandparent, up to the root. Those are its descriptors — the categories the system uses to characterize it.

Nested categories are parallel categories. A set and its superset do not point in different directions; one is contained in the other, and the smaller carries no information the larger contradicts. Under the encoding this manual uses, the pair contributes u ~ 0.

We badge the step: **[POSTULATE] a pair of nested categories is assigned u = 0, and a pair of incomparable (overlapping or disjoint but non-nested) categories is assigned some u_0 > 0 determined by measurement.** The postulate is a modeling choice; it is defensible because a nested pair conveys strictly refined information along a single axis, and it is testable by measuring u between a category and its parent using Chapter 7's instruments.

With that postulate, the arithmetic is stark. Consider an element carrying m descriptors:

TABLE 8.10 — DESCRIPTOR-SET AUDIT FOR ONE ELEMENT CARRYING m LABELS
(semilattice case computed at measured u_0 = 0.7)

m	TREE: all nested		SEMILATTICE: all incomparable		ceiling
	N_eff	S	N_eff	S	ln m
2	1.000	0.0000	1.700	0.5306	0.6931
3	1.000	0.0000	2.400	0.8755	1.0986
4	1.000	0.0000	3.100	1.1314	1.3863
6	1.000	0.0000	4.500	1.5041	1.7918
8	1.000	0.0000	5.900	1.7750	2.0794
12	1.000	0.0000	8.700	2.1633	2.4849

In a strict tree, an element's entire descriptor set audits to an effective count of one, regardless of how deep the hierarchy is. Twelve levels of nested classification and the element is described by one direction. Add levels and you add labels, cost, and review steps, and you add exactly zero surplus.

Real organizations are not perfect trees and not complete semilattices. Parameterize by the fraction q of an element's descriptor pairs that are incomparable rather than nested:

TABLE 8.11 — PARTIAL OVERLAP. m = 6 DESCRIPTORS, u_0 = 0.7

q (fraction incomparable)	u_bar	N_eff	S	% of ln 6
0.0	0.000	1.000	0.0000	0.0%
0.2	0.140	1.700	0.5306	29.6%
0.4	0.280	2.400	0.8755	48.9%

0.6	0.420	3.100	1.1314	63.1%
0.8	0.560	3.800	1.3350	74.5%
1.0	0.700	4.500	1.5041	83.9%

Moving an organization from q = 0.0 to q = 0.2 — allowing one descriptor pair in five to overlap rather than nest — buys 29.6% of the available ceiling for that element's descriptor set, and it is the cheapest structural change in this chapter.

What does q > 0 look like in practice? It looks like:

- A person who reports to a function and also belongs to a standing cross-function body, where neither membership is a subset of the other.
- A budget line owned jointly by two units, neither of which contains the other.
- A physical space used by two groups on overlapping schedules — Jacobs' mixed primary uses, which is the same structure in a city.
- A document with two independent owners in different branches.
- A review process in which a proposal is read by its own division and by a body that spans divisions, with neither review subordinate to the other.

And what does q = 0 look like? It looks like a clean org chart. The org chart is legible precisely because it is a tree, and it is a tree precisely because someone removed all the overlaps, and the overlaps were the surplus. This is Alexander's complaint about planned cities restated for institutions, and it is the same complaint: the tree is what a designer produces when the design must be comprehensible to the designer, and comprehensibility to the designer is not one of the system's requirements.

Chapter 20 develops the organizational consequence (silos as local minima). Chapter 12 develops the spatial consequence. Here the point is arithmetic: **hierarchy is a u-destroying structure, and since surplus is `ln(1 + (N−1)u)`, destroying u destroys the return on every category you have already paid for.**

The semilattice is not free Honesty requires the other side. Overlapping structure raises C. A person with two non-nested memberships attends two sets of meetings, answers to two review paths, and can be given contradictory instructions. Alexander's essay does not price this; this manual must.

The cost of overlap is roughly proportional to the number of overlapping memberships per element, which is a linear term, while the surplus gain is inside a logarithm. Therefore overlap, like every other lever, saturates: the first overlap per element is worth a great deal and the sixth is worth very little, while each costs the same. Table 8.11 shows the shape — most of the gain is in the first fifth of q. The design rule follows in the Specification: give every element one or two non-nested memberships and stop.

8.11 The four legitimate reasons to raise N

Given the prohibition, when is raising N correct?

L-1 — Ashby shortfall. The environment presents distinguishable variety the system cannot distinguish. Diagnostic signature: recurring surprises of a *kind*, not just of a magnitude — failures the system has no category for and therefore no procedure for. This is the strongest reason and it is not optional; a system without requisite variety fails at its function regardless of its headroom.

Note carefully: this is a capability argument, and it may require accepting a worse headroom position in exchange for being able to do the job at all. That trade is legitimate and must be made explicitly, with the headroom cost computed and stated.

L-2 — Building, not repairing. At the initial design of a system, N is being set for the first time and the concavity of Table 8.2 means the early categories are cheap and valuable. Getting from one category to four buys 1.3863 nats. Do this at design time, when C is still small and Q is still zero, and never again at this price.

L-3 — The ceiling is binding and h is positive. If η S is running up against η ln N — that is, if u is already near 1 and the system is genuinely extracting near its ceiling — then N is the binding constraint and raising it is the correct move. This condition is rare and must be measured, not assumed. The test: `S / ln N > 0.85`. Below that, u is the binding constraint, not N, and Chapter 9 is the correct chapter.

L-4 — The new category is nearly free. Occasionally a genuinely distinct category can be acquired at near-zero marginal C — an existing person who already holds a second practice, an existing space that already serves a second use, an existing artifact that already encodes a second tradition. **Take these every time.** They are the only case where the logarithm's slowness does not matter, because the denominator is zero.

Everything else is category inflation.

8.12 N against cost — the crossover

Chapter 5 established the crossover calculation; it is reproduced here in the form that governs N decisions, because it is the arithmetic that turns the prohibition from an argument into a number.

Take a coordination cost that grows with communication paths: `C(n) = k . n(n−1)/2`. Compare against η ln n at η = 0.5.

TABLE 8.12 — COST VS CEILING, C(n) = k n(n−1)/2 AGAINST eta ln n, eta = 0.5

n	k=0.002		k=0.005		k=0.01		k=0.02		k=0.05	
	C	net	C	net	C	net	C	net	C	net
4	0.012	+0.681	0.030	+0.663	0.060	+0.633	0.120	+0.573	0.300	+0.393
6	0.030	+0.866	0.075	+0.821	0.150	+0.746	0.300	+0.596	0.750	+0.146
8	0.056	+0.984	0.140	+0.900	0.280	+0.760	0.560	+0.480	1.400	−0.360
10	0.090	+1.061	0.225	+0.926	0.450	+0.701	0.900	+0.251	2.250	−1.099
12	0.132	+1.111	0.330	+0.913	0.660	+0.583	1.320	−0.078	3.300	−2.058
15	0.210	+1.144	0.525	+0.829	1.050	+0.304	2.100	−0.746	5.250	−3.896
20	0.380	+1.118	0.950	+0.548	1.900	−0.402	3.800	−2.302	9.500	−8.002
25	0.600	+1.009	1.500	+0.109	3.000	−1.391	6.000	−4.391	15.000	−13.391

```
     30 | 0.870   +0.831    | 2.175   −0.474    | 4.350   −2.649    | 8.700   −
6.999    |21.750 −20.049
     40 | 1.560   +0.284    | 3.900   −2.056    | 7.800   −5.956    |15.600 −
13.756   |39.000 −37.156
```

Every column turns negative. Every one. The only question is where. At k = 0.01 — a modest coordination cost of one hundredth of a nat per communication path — the system goes net-negative at n = 20 and is 5.96 nats underwater at n = 40, against a ceiling of only 1.84.

Note also where each column peaks: k = 0.002 peaks around n = 15; k = 0.005 around n = 10; k = 0.01 around n = 8; k = 0.02 around n = 6; k = 0.05 around n = 4. **The optimal category count is small, and it shrinks fast as coordination gets expensive.** This is the arithmetic behind every piece of practical folk wisdom about small teams, and it does not require any claim about human psychology. It is a logarithm losing a race against a quadratic.

The Specification

RN-1 — Publish the ceiling. Every system's design documentation must state `N`, `ln N`, η, and `η ln N` on its first page. That last number is the total budget from which all maintenance, coordination, and growth is paid, and if the people spending it do not know its value they are spending blind.

RN-2 — Compute h before considering N. `h = η ln N − δQ − C`. If h < 0, the Domain-Addition Prohibition applies and no proposal to raise N may be approved as a remedy. Route to Chapter 26's decision procedure: cut C, raise η, or shed Q.

RN-3 — Run the relabeling audit on every proposed category. Six steps, §8.7. A proposal that cannot state the new category's encoder in one sentence is rejected without further review.

RN-4 — Require u >= 0.3 against every existing category. Measured by Chapter 7's Instrument 1 or 2, on at least 20 discriminating items, under blind elicitation. Report the full vector of measured u's against existing categories, not the mean.

RN-5 — Price ceiling gain and delivered gain separately. `ln(N+1) − ln N` is the ceiling gain and is nearly always trivially small (Table 8.7). The delivered gain is `ln(1 + N.u_bar_new) − ln(1 + (N−1).u_bar_old)` and it is the number that matters. Documents that quote only the ceiling gain are non-compliant.

RN-6 — Test the binding constraint before acting. Compute `S / ln N`. Above 0.85, N binds and this chapter applies. Below 0.85, u binds and Chapter 9 applies. Most systems that believe they have an N problem have a u problem, and this single ratio distinguishes them.

RN-7 — Diagnose Ashby shortfall separately from headroom deficit. Ashby shortfall presents as surprises of a novel *kind*. Headroom deficit presents as steady erosion of a familiar kind. Keep two registers: a surprise log (novel failure kinds, dated) and a headroom trace (h computed each period). They call for opposite actions and must never be conflated.

RN-8 — Specify q, the overlap fraction, as a design parameter. For each element of the organization, count its memberships and the fraction of membership pairs that are non-nested. Target q in the range 0.2 to 0.4 (Table 8.11 shows the gain is concentrated there and Chapter 5's cost accounting shows the overhead is not yet severe). Publish q per element. q = 0 across the

whole organization means you have built a tree and forfeited the return on every category you paid for.

RN-9 — Give every person at least one non-nested membership. The minimum viable semi-lattice. One standing cross-function body, one jointly owned artifact, or one shared space with a group that neither contains nor is contained by their own.

RN-10 — Cap N against the cost crossover. Compute the crossover from Table 8.12's method at your measured k and η. Do not exceed the peak. Publish the computed optimum with the org design.

RN-11 — Prefer splitting the body to growing it. Two bodies of six at u = 0.5 deliver `2 x 1.5041 = 3.0082` nats against `2 x 15 = 30` paths. One body of twelve at the same u delivers 1.8718 against 66 paths. **Splitting raises delivered surplus by 61% and cuts communication paths by 55% simultaneously.** This is the strongest single move available at moderate N and it is almost never taken, because splitting reads as failure and growing reads as success.

RN-12 — Badge every use of the Ashby correspondence. `[FORMAL CORRESPONDENCE, NOT IDENTIFICATION]`. Ashby's variety and the surplus ceiling are both logarithms of a state count; they are not proved equal and Paper D declines to identify surplus with information-theoretic quantities.

RN-13 — Badge the nesting postulate. Any argument that assigns u = 0 to a nested category pair carries `[POSTULATE]` and states the test that would refute it: measure u between a category and its parent using Instrument 2; if it exceeds 0.3, the postulate does not hold for that structure.

The Failure Modes

FM-8.1 . THE DOMAIN-ADDITION REFLEX

Damages: C (raises it), h (worsens it), and the credibility of the whole planning process.

Clinical description. The system detects decline. Leadership responds by broadening scope: new departments, new initiatives, new disciplines, new markets. The reasoning is that a bigger system has more capability and therefore more resilience. The arithmetic says the added ceiling is `η ln(N_new/N_old)` while the added cost is at least linear and typically quadratic in the new category count. **Every domain addition made in response to a deficit deepens the deficit.**

Diagnostic signature. A strategy document that responds to declining performance with an expanded remit. Growth in the number of named initiatives while output per initiative falls. The phrase "we need to be more than just a ___."

Detecting measurement. Compute h before and after each of the last three expansions. If h fell each time, the reflex is present and is the primary pathology.

Fix. RN-2 as a hard gate. Then Chapter 26's ordering: raise η first (cheapest), cut C second, shed Q third. Only after h > 0 may N be reconsidered, and then only under one of §8.11's four legitimate reasons.

FM-8.2 . CATEGORY INFLATION

Damages: C, and the integrity of the measurement of N.

Clinical description. N in the plan rises while N in the geometry does not. Categories are created by splitting, relabeling, or nominal reassignment. Table 8.7's ceiling gain is quoted; Table 8.7's ceiling gain is 1.63% at N = 20 and the delivered gain is zero.

Diagnostic signature. Growth in the org chart without growth in scored disagreement. New category names that are compound modifications of existing ones ("Applied X", "Strategic X", "X Operations"). Measured u between new and parent category below 0.15.

Detecting measurement. RN-4. Blind elicitation on 20 discriminating items, new category against its parent, computed as $u = 1 - \rho^2$.

Fix. The relabeling audit (§8.7) applied retroactively. Merge every category pair measuring $u < 0.15$ back into one. Report the merge as a cost reduction, which is what it is.

FM-8.3 . THE ASHBY–HEADROOM CONFUSION

Damages: the diagnosis, and therefore every action that follows it.

Clinical description. A capability failure (missing variety) and a persistence failure (negative headroom) present superficially alike — both look like "we are not doing well" — and call for opposite actions. Treating a headroom deficit as an Ashby shortfall produces FM-8.1. Treating an Ashby shortfall as a headroom deficit produces a system that is efficiently and sustainably unable to do its job.

Diagnostic signature. For Ashby shortfall: failures of a *novel kind*, incident reports whose category field reads "other," repeated invention of ad-hoc procedure. For headroom deficit: failures of a familiar kind at rising frequency, maintenance backlogs, staff who know exactly what is wrong and cannot get to it.

Detecting measurement. Two registers, kept separately, per RN-7: a surprise log and a headroom trace.

Fix. Diagnose first. If Ashby shortfall: acquire the missing distinction, accept the headroom cost, state it explicitly, and pay for it by cutting elsewhere. If headroom deficit: Chapter 26's procedure. If both — which happens — fix headroom first, because an Ashby intervention undertaken at $h < 0$ will not persist. You cannot store a capability in a vessel that is draining.

FM-8.4 . THE TREE

Damages: u, catastrophically, and therefore the return on all existing N.

Clinical description. The organization has been arranged as a strict hierarchy. Every element has exactly one parent. Every category is a subset of exactly one other. All overlaps have been resolved into ownership, all dual reporting eliminated, all shared spaces assigned. The chart is clean. Under the nesting postulate, every element's descriptor set audits to $N_{eff} = 1$.

Diagnostic signature. q = 0 across the organization. No standing cross-function bodies with real authority. Every artifact has exactly one owner. Every meeting has exactly one convening

unit. When two units need to work together, a coordinating third unit is created above them — which preserves the tree and adds a level.

Detecting measurement. RN-8. Enumerate memberships per element and compute the non-nested fraction. Also count: how many standing bodies span two branches without being subordinate to a common parent below the root? In a strict tree the answer is zero.

Fix. RN-9, the minimum viable semilattice: one non-nested membership per person. Then targeted overlap at the highest-value boundaries — Chapter 20's rotation programs, shared instrumentation, and boundary objects are all q-raising devices. Do not attempt $q > 0.4$; Table 8.11 shows the returns saturate and Chapter 5 shows the costs do not.

FM-8.5 . THE UNMEASURED CEILING

Damages: every planning decision downstream.

Clinical description. Nobody has computed $\eta \ln N$. Budgets, growth targets, and maintenance commitments are set without reference to the total from which they must all be paid. The system routinely commits to $r_t > h_t$ and then fails to deliver, and the failure is attributed to execution.

Diagnostic signature. Chronic under-delivery against plan with no pattern in which items fail. Plans that sum commitments without reference to a total. Nobody in the organization can state $\eta \ln N$ when asked.

Detecting measurement. Ask three senior people for the value of $\eta \ln N$. If none can answer, the mode is present.

Fix. RN-1. Put the number on the first page of the plan. This costs one afternoon and it is, per unit of effort, the highest-return intervention in Part II.

FM-8.6 . THE SATURATED CATEGORY SET

Damages: C, while delivering nothing to S.

Clinical description. N has grown past the crossover in Table 8.12. Additional categories cost more in coordination than they return in ceiling. The system is net-negative on its most recent several categories and does not know it because it measures ceiling gain rather than delivered gain minus cost.

Diagnostic signature. Coordination effort growing faster than headcount. Meetings whose purpose is to synchronize other meetings. Category owners who cannot name what their category contributes that no other category does.

Detecting measurement. Compute the crossover per RN-10 at measured k and η. Compare against current N.

Fix. RN-11. Split rather than grow, or merge the categories that fail RN-4. Both reduce the quadratic term, and the surplus loss is small because the categories being removed were the ones contributing 0.02–0.05 nats apiece.

Worked Example

The research institute that wanted more domains

A publicly funded research institute has operated for nineteen years. It spans eight declared research domains. Its stated problem, from its own strategic review: "our output per researcher has declined for six consecutive years and our field-leading position has eroded." Its proposed remedy: expand to sixteen domains over five years, on the reasoning that a broader institute is a more resilient one and that the adjacent fields are where the growth is.

We survey before touching anything (Chapter 24's procedure, abbreviated here to the parameters this chapter needs).

Step 1 — Parameters

Parameter	Value	How obtained
N	8	Declared research domains, each with its own funding line, faculty, and review process. Relabeling audit run: all eight pass RN-4 with pairwise u >= 0.42.
η	0.40	Chapter 4's method: attributable ΔQ per unit S over the last six years. Mid-range; the institute publishes and archives but converts little of its internal interaction into durable method, training, or instrumentation.
δ	0.06 per year	Chapter 3's method: decay of the order proxy under a two-year period of reduced input following a funding pause in year 12.
Q	40.0	Declared order proxy: a composite of maintained instrument-hours, active protocol count, and trained-researcher count, normalized to the year-1 baseline.
C	0.55 per year	Chapter 5's tally: governance, compliance, inter-domain coordination, estate, and the review overhead of eight funding lines.

Step 2 — The ceiling and the gauge

```
ln N       = ln 8              = 2.0794
eta ln N   = 0.40 x 2.0794     = 0.8318      <- the ENTIRE annual budget

delta Q    = 0.06 x 40.0       = 2.4000
C          =                     0.5500
-------------------------------------------
h = eta ln N - delta Q - C
  = 0.8318 - 2.4000 - 0.5500
  = -2.1182
```

The institute's entire annual budget of order-production is 0.8318 nats, and its maintenance burden alone is 2.4000. Headroom is -2.1182. By Corollary 3.4 this is the regime in which **no choice of interaction geometry prevents order from decreasing.** Six consecutive years of decline is not a mystery, a morale problem, or a leadership failure. It is the equation.

Note what this means about the six years of effort already spent. Whatever the institute did to improve its collaboration, its culture, its recruitment, or its geometry, S is bounded above by ln 8 = 2.0794 and ηS by 0.8318. The institute could have achieved perfect orthogonality across all eight domains — u = 1, every pair — and still declined by 2.1182 - 0 = 1.5682 nats per year at best. **The geometry was never the binding constraint and no geometric intervention could have worked.**

Step 3 — Price the proposed remedy Expand from 8 domains to 16.

```
ln 16      = 2.7726
eta ln 16  = 0.40 x 2.7726 = 1.1090
gain       = 1.1090 - 0.8318 = 0.2772 nats
```

Doubling the institute's entire domain structure — five years of recruitment, eight new funding lines, eight new estates, eight new review processes — buys **0.2772 nats** against a deficit of 2.1182. It closes 13% of the gap. And that is before counting the cost.

Continue the sequence to make the point unanswerable:

TABLE 8.13 — CLOSING A 2.1182-NAT DEFICIT BY DOMAIN ADDITION ALONE
(eta = 0.40, delta Q = 2.4000, C = 0.5500 held constant)

N	eta ln N	h = eta ln N - 2.9500
8	0.8318	-2.1182
16	1.1090	-1.8410
32	1.3863	-1.5637
64	1.6636	-1.2864
128	1.9408	-1.0092
1000	2.7631	-0.1869
1600	2.9511	+0.0011

Solve it exactly. Break-even requires η ln N >= δQ + C = 2.9500, hence ln N >= 2.9500 / 0.40 = 7.3750, hence

```
N  >=  e^{7.3750}  =  1595.6
```

The institute would need 1,596 genuinely distinct research domains to reach break-even by domain addition alone — each one passing RN-4's u >= 0.3 test against all 1,595 others, each with its own staff, and all of them held together at the institute's current coordination cost, which they obviously would not be. With 1,596 domains the coordination term alone, at any positive k, is astronomically larger than 0.55.

This is the Domain-Addition Prohibition in a single number. Print it and put it on the wall.

Step 4 — Price the admissible levers *Lever 1 — Raise η.* The institute's η = 0.40 reflects a real and fixable deficiency: internal seminars produce no artifact, methods live in individual heads, instrument configurations are not documented, and departing postdocs take their protocols with them. Chapter 4's interventions — mandatory method write-ups, a versioned protocol archive, apprenticeship overlap on instrument handover — are all cheap.

```
eta = 0.40:   h = 0.40 x 2.0794 - 2.95 = -2.1182
eta = 0.60:   h = 0.60 x 2.0794 - 2.95 = -1.7023
eta = 0.90:   h = 0.90 x 2.0794 - 2.95 = -1.0785
eta = 1.20:   h = 1.20 x 2.0794 - 2.95 = -0.4547
eta = 1.42:   h = 1.42 x 2.0794 - 2.95 = +0.0028   <- break-even
```

η would have to more than triple, to 1.42, to close the deficit alone. Tripling η is ambitious but it is not absurd — an institute that currently captures almost nothing in durable form has a great deal of room — and unlike the 1,596 domains it is at least the right order of magnitude.

Lever 2 — Cut C.

```
C = 0.55:  h = 0.8318 - 2.40 - 0.55 = -2.1182
C = 0.30:  h = 0.8318 - 2.40 - 0.30 = -1.8682
C = 0.10:  h = 0.8318 - 2.40 - 0.10 = -1.6682
C = 0.00:  h = 0.8318 - 2.40 - 0.00 = -1.5682
```

Even eliminating coordination cost entirely — an impossibility — leaves h at -1.5682. C is not the problem. The δQ term is.

Lever 3 — Shed Q.

```
Q = 40.0:  h = 0.8318 - 2.400 - 0.55 = -2.1182
Q = 30.0:  h = 0.8318 - 1.800 - 0.55 = -1.5182
Q = 20.0:  h = 0.8318 - 1.200 - 0.55 = -0.9182
Q = 10.0:  h = 0.8318 - 0.600 - 0.55 = -0.3182
Q =  4.7:  h = 0.8318 - 0.282 - 0.55 = -0.0002   <- break-even
```

The institute would have to shed 88% of its stored order to break even on this lever alone. That is the arithmetic statement of a nineteen-year-old institution that has accumulated more structure than its input can maintain. This is the maintenance trap of Chapter 6, and it is the actual diagnosis: the institute is not failing to produce; it is failing to stop maintaining.

Step 5 — The combination that works No single lever suffices. Combine, at levels that are individually plausible:

```
eta:  0.40 -> 0.60  (method archive, protocol versioning, handover overlap)
 C:    0.55 -> 0.25   (merge four of eight funding lines into two; single
                  compliance function; retire two of five governance bodies)
Q:   40.0 -> 18.0  (retire 55% of stored order: decommission the instrument
                 park's obsolete half, sunset protocols unused in four years,
                  close two domains that fail RN-4 against their neighbours)

 h = 0.60 x 2.0794 - 0.06 x 18.0 - 0.25
   = 1.2476 - 1.0800 - 0.2500
   = -0.0823
```

Still marginally negative. Push η slightly further, to 0.68, which the protocol archive plus a formal apprenticeship overlap should support:

```
h = 0.68 x 2.0794 - 1.0800 - 0.2500 = 1.4140 - 1.3300 = +0.0840
```

Positive. Barely — but the sign is what matters, because the sign is the difference between a system that can grow and a system for which no policy exists that prevents decline.

Note what happened to N. Two domains were **closed**, taking N from 8 to 6:

```
eta ln 6 = 0.68 x 1.7918 = 1.2184
h = 1.2184 - 1.0800 - 0.2500 = -0.1116
```

Negative again. The closures cost 0.1956 nats of ceiling. So they must be paid for by the C and Q reductions they enable — which they are, and generously: closing two domains was what made the funding-line merge and the instrument decommissioning possible in the first place. Attribute honestly:

```
    closing 2 domains:  ceiling  -0.1956
                           delta Q  -0.36    (from Q 24 -> 18, the domain-
attributable part)
                        C        -0.12
                        net h    +0.2844
```

Closing two domains improved headroom by 0.2844 nats, while adding eight domains would have improved it by at most 0.2772 before costs and by a substantial negative after them. That is the chapter's thesis reduced to a single comparison, and it is the comparison no strategic review ever makes.

Final position:

```
N = 6,  eta = 0.68,  delta = 0.06,  Q = 18.0,  C = 0.25
eta ln N = 0.68 x 1.7918 = 1.2184
h        = 1.2184 - 1.0800 - 0.2500 = -0.1116
```

which is still negative, so one more cut is required. Take Q from 18.0 to 15.0 by retiring the third governance body and the associated documentation estate:

```
h = 1.2184 - 0.9000 - 0.2500 = +0.0684
```

Positive. The institute is smaller in every dimension and it is, for the first time in seven years, above water.

Step 6 — Now check Ashby (RN-7) Before accepting the recommendation, run the separate diagnosis. The institute's surprise log over six years shows 31 incidents. Classifying by kind:

- 24 are familiar failures at rising frequency: instrument downtime, grant deadline misses, staff departures without handover. **Headroom signature.**
- 7 are novel in kind: three regulatory categories the institute had no procedure for, two data-governance failures involving a technique nobody on staff could evaluate, two collaborations that failed because no one could translate between two of the eight domains. **Ashby signature.**

The seven are real and they will not be fixed by the headroom intervention. Two of them — the data-governance failures — indicate a genuine missing distinction, which is L-1, the one legitimate reason that overrides headroom concerns.

So the recommendation carries an explicit exception:

```
Acquire ONE new category: quantitative data governance, staffed by one person
 with genuinely distinct training (RN-4 test required before hiring).

ceiling gain:  eta [ln 7 - ln 6] = 0.68 x (1.9459 - 1.7918) = 0.68 x 0.1541 = +0.1048
 C cost:         estimated +0.06 (one funding line, one review path)
 Q cost:         +0 initially
 net h:          +0.0684 + 0.1048 - 0.06 = +0.1132
```

In this case the Ashby acquisition happens to be **h-positive**, because the institute is at N = 6 where Table 8.7's ceiling gain per category is still 8.60% of the current ceiling, and because the new category is cheap. That is L-4 and L-1 arriving together. **Take it.** Had it been h-negative, the correct move would have been to take it anyway and state the cost, because a system that cannot distinguish its regulatory environment fails at its function regardless of its budget — but the cost would have had to be paid from somewhere, explicitly, in the same document.

Step 7 — What the plan says

```
RECOMMENDATION — RESEARCH INSTITUTE, YEAR 20

REJECT: expansion from 8 to 16 domains.
  Reason: h = -2.1182. Domain-Addition Prohibition (Chapter 8 §8.6).
  Break-even by domain addition alone requires N = 1,596.
  Proposed expansion closes 13% of the deficit before costs and
  widens it after them.

ADOPT, in this order:
  1. RAISE eta 0.40 -> 0.68. Method archive, protocol versioning,
     mandatory handover overlap. Measured per Chapter 4.
  2. CUT C 0.55 -> 0.25. Merge funding lines 8 -> 4; single compliance
     function; retire 3 of 5 governance bodies.
  3. SHED Q 40.0 -> 15.0. Decommission obsolete instrument park; sunset
     four-year-unused protocols; close 2 domains failing RN-4.
  4. THEN acquire 1 category (data governance) on Ashby grounds L-1,
     at measured u >= 0.3 against all existing domains.
```

```
RESULTING POSITION
  N = 7,  eta = 0.68,  delta = 0.06,  Q = 15.0,  C = 0.31
  eta ln N = 0.68 x 1.9459 = 1.3232
  h        = 1.3232 - 0.9000 - 0.3100 = +0.1132
  Q_inf    = (eta F(u*) - C) / delta, computed once u* is measured
             post-intervention (Chapter 9).
```

Step 8 — Falsification (Chapter 25)

- **If h computed at year 21 has not risen above zero**, at least one of the three parameter changes did not land, and the survey must identify which. The most likely candidate is the Q shed: decommissioning is routinely announced and not executed, and Q is the parameter most prone to being reported as reduced while still being maintained.
- **If output continues to decline while h > 0**, the order proxy is wrong. Paper E §6.3's warning applies: verify that the declared proxy Y is not merely a repackaging of the input U. This is the most common measurement error in the framework and Chapter 24 §24.7 gives the test.
- **If the shuffled-input control (Chapter 25) fits as well as the ordered one**, the apparent relationship between the institute's interaction structure and its stored order is an artifact and the entire analysis must be withdrawn. This is the control most often omitted and it is the one that would kill this analysis if it were wrong.
- **If the seven Ashby-signature surprises continue at the same rate after the data-governance acquisition**, the missing distinction was misidentified. Re-classify the surprise log and look for the actual missing category rather than assuming the first candidate was correct.

The ceiling is a logarithm and you cannot outrun it. That is the negative result. The positive one is waiting in the term the ceiling multiplies: `S = ln(1 + (N−1)u)`, in which u enters linearly and multiplies the full `(N−1)`, and in which u — unlike N, unlike δ, unlike Q — can be moved in a quarter.

That is the next chapter, and it is the one to read if you have a system to fix this year.

Chapter 9 . Engineering u — Designing Non-Parallelism

Chapter 8 closed a door. The ceiling is `ln N`, it grows brutally slowly, and a headroom deficit cannot be closed by adding domains. That is the negative result and it eliminates the move most organizations reach for first.

This chapter opens the door that remains. In `S = ln(1 + (N−1)u)`, the parameter u enters **linearly**, and it multiplies the **full** `(N−1)`. Every category you have already paid for — every domain, every discipline, every hire, every instrument — is sitting there with its contribution scaled by u. Raise u and you collect on all of them at once, without hiring anyone, without expanding scope, and without adding a single communication path.

u is the only parameter in the law that can be moved substantially in a single quarter, and it is the only one whose improvement collects on capital you have already bought. δ moves over years. η moves over seasons and requires building things. N is capped by a logarithm and constrained by cost. C can be cut but only to zero, and zero is rarely enough. u can go from 0.12 to 0.50 in eight quarters with three interventions and no new headcount, and this chapter shows exactly how, with the arithmetic attached.

It also catalogs, at length, the many ways institutions destroy u without noticing — because in practice most u engineering is not construction. It is the removal of homogenizing machinery that somebody installed for a good reason and nobody has audited since.

The Law

9.1 The parameter

```
u(x,y) = 1 − (x . y)^2 = || x ^ y ||^2  = sin^2 θ        in [0,1]

N_eff  = 1 + (N−1) u

S      = f(u) = ln(1 + (N−1) u)      UNIQUELY characterized by S1–S4
```

with **S3, the affine effective count, being the design axiom — chosen, not derived**. The linearity of `N_eff` in u, which is the entire engineering premise of this chapter, is a consequence of S3. Say so whenever you lean on it, which in this chapter is constantly.

The derived properties that govern u engineering (Paper A):

Property	Ref	What it means for the engineer
Strict monotonicity	Thm 3.1	Every increase in u strictly increases S. There is no threshold to cross and no plateau to wait for. **Partial progress is real progress.**

Property	Ref	What it means for the engineer
Positivity	Thm 3.2	`F > 0` for any `u > 0`. A single genuinely different member produces strictly positive surplus.
Strict concavity	Thm 3.3	Diminishing returns. The first increments of u are the valuable ones.
Sharp Lipschitz	Thm 4.2	`|f(u_1)−f(u_2)| <= (N−1)|u_1−u_2|`. Sensitivity is highest at small u, which is where most systems are.
Extremality	Thm 4.3	`F = 0` iff parallel. **Not small. Zero.**
Range	Paper D Cor 2.2	`0 <= F <= ln N`.

9.2 The zero

Theorem 4.3 says `F = 0` if and only if the states are parallel. This is not a limit or an approximation. At u = 0, `N_eff = 1` and `S = ln 1 = 0` exactly.

The organizational statement of that theorem is the sentence this chapter is built around:

> **A second opinion from someone whose training and incentives replicate the first is worth ln(1) = 0.**

Not "worth little." Not "of limited value." Worth zero. The second reviewer, the confirming analysis, the sanity read from the colleague down the hall — if the instrument is the same instrument, the reading is the same reading, and the law prices a duplicate at nothing.

TABLE 9.1 — WHAT A SECOND OPINION IS WORTH, AS A FUNCTION OF ERROR CORRELATION
(N = 2; u = 1 − rho^2; S = ln(1 + u); ceiling ln 2 = 0.6931)

rho	u	N_eff	S = ln(1+u)	share of ln 2
1.00	0.0000	1.0000	0.0000	0.0%
0.98	0.0396	1.0396	0.0388	5.6%
0.95	0.0975	1.0975	0.0930	13.4%
0.90	0.1900	1.1900	0.1740	25.1%
0.80	0.3600	1.3600	0.3075	44.4%
0.60	0.6400	1.6400	0.4947	71.4%
0.40	0.8400	1.8400	0.6098	88.0%
0.20	0.9600	1.9600	0.6729	97.1%
0.00	1.0000	2.0000	0.6931	100.0%

Read the top of the table, because the top is where organizations live. Two reviewers at ρ = 0.95 — typical of two people trained in the same programme, reading the same literature, evaluated by the same committee — deliver 13.4% of an independent pair. At ρ = 0.98, which describes a mentor and their own former student, they deliver 5.6%.

The organization believes it bought a second opinion. It bought 5.6% of one, at full price.

And now the other end of the table, which is not printed above and must be. Table 9.1 runs `rho` from 1.00 down to 0.00 and stops, which invites the reading that the good direction is simply *away from agreement*, and that the best available reviewer is the one who disagrees hardest. That reading is wrong, and the correction is one line of the definition: `u = 1 − rho^2` depends only on `|rho|`. The table is symmetric about zero.

```
rho = −0.95  ->  u = 0.0975  S = 0.0930   13.4%  -- identical to rho = +0.95
rho = −0.98  ->  u = 0.0396  S = 0.0388    5.6%  -- identical to rho = +0.98
rho = −1.00  ->  u = 0.0000  S = 0.0000    0.0%  -- identical to rho = +1.00
```

At perfect opposition the surplus is exactly zero, precisely as it is at perfect agreement. The zero of Theorem 4.3 is a zero of *parallelism*, and a state vector and its negation are parallel. Given the sycophant's reading you can reconstruct the principal's; given the perfect contrarian's you can reconstruct it just as exactly, by changing a sign. **The mirror and the mirror-image are the same instrument. The contrarian who inverts you carries as much information as the sycophant who echoes you: none.**

This correction governs the rest of the chapter, and method M-7 and the red-team specification in particular. **The target of every intervention below is `|rho|` near zero — genuine independence — and never `rho` near -1.** A body engineered for "creative tension," in which a designated opponent takes the negation of whatever is proposed, has manufactured a u of zero and is paying `C` for the extra seat: it returns no surplus and costs more than the aligned pair it replaced. Chapter 7 §7.6 gives the geometry — the angle table folded about 90 degrees — and Chapter 11 catalogues the fault in its own right, distinct from the Parallel Wall.

9.3 The honesty clauses, restated for this chapter

1. **Paper A is pure mathematics and makes no empirical or metaphysical claim.** That `F = 0` at u = 0 is a theorem about a functional form. That your review panel's second reader is worth nothing is a modeling inference, and it depends on the encoder you chose and on the Roster Reduction [POSTULATE] of Chapter 7 §7.4.
2. **S3 is a design axiom.** The linearity of `N_eff` in u is chosen.
3. **The deployed functional is a lower bound.** A measured u of 0.15 licenses "at least this much"; it does not license "exactly this much" or "no more than this."
4. **The identity `u = 1 − ρ²` is exact** for mean-centred, unit-normalized state vectors, since their inner product is the Pearson correlation. Encoding a member's state as their error vector is a modeling assumption, badged **[POSTULATE]**.
5. **Subadditivity holds on u only; the triangle inequality fails for geometric triples** (Paper A §5). When this chapter combines interventions, it combines them on u and does not chain surplus across triples.
6. **Surplus is not Shannon entropy or mutual information** (Paper D), and **no thermodynamic negentropy law is proved** (Paper E).

9.4 The lever ranking, and why u is second

Chapter 26 gives the full ordering. The short version, with this chapter's justification:

Lever	Typical time to move	Typical cost	Collects on	Verdict
η	one to three quarters	low — writing, archiving, apprenticeship	all future S	**first move in almost every system**
u	one to four quarters	low to moderate — mostly removal of existing machinery	all existing N	**fastest-moving parameter; this chapter**
C	one to four quarters	moderate — political, not technical	the whole budget	third; bounded below by zero
Q (shedding)	one to eight quarters	political	frees δQ	underused; see Chapter 6
δ	years	high	future retention	slow; Chapter 3
N	years	very high	logarithmic return	**last; Chapter 8's prohibition**

η is first because it is usually cheapest and because it multiplies everything. But η requires **building** — an archive, a protocol, a training path — and building takes a quarter before it starts converting. u is the parameter that moves fastest, because most u interventions are subtractions rather than additions, and a subtraction takes effect the day it is made.

Cancel the meeting that produces the cascade and u rises the next week. Stop screening for culture fit and u rises with the next hire. Remove the shared metric that gives every unit the same objective function and u rises the next quarter. These are not construction projects. They are the removal of machinery.

The Engineering

9.5 What u buys, at a working scale

Before methods, the payoff table. A system with N = 10, η = 0.5, δ = 0.08 per quarter, C = 0.30. Each level of u, with the steady state it implies through Theorem 3.5, `Q_inf = (ηF(u*) − c*)/δ`.

`TABLE 9.2 — WHAT u BUYS, N = 10, eta = 0.5, delta = 0.08, C = 0.30`

u	N_eff	S = ln N_eff	eta S	Q_inf = (eta S − C)/delta
0.05	1.45	0.3716	0.1858	−1.428 (unsustainable)
0.10	1.90	0.6419	0.3209	0.262
0.15	2.35	0.8544	0.4272	1.590
0.20	2.80	1.0296	0.5148	2.685
0.30	3.70	1.3083	0.6542	4.427
0.40	4.60	1.5261	0.7630	5.788

0.50	5.50	1.7047	0.8524	6.905
0.60	6.40	1.8563	0.9281	7.852
0.70	7.30	1.9879	0.9939	8.674
0.80	8.20	2.1041	1.0521	9.401
0.90	9.10	2.2083	1.1041	10.052
1.00	10.00	2.3026	1.1513	10.641

Three readings.

First, the bottom of the table is a death sentence. At u = 0.05 the steady state is negative, which by Theorem 3.5 means the equilibrium is unsustainable and the linear model must be replaced. The theorem does not hide this failure mode and neither should you. A ten-category system at u = 0.05 does not converge to a little stored order; it converges to a contradiction, which in practice means it dissolves, is rescued by subsidy, or was never described by these parameters.

Second, the returns are front-loaded. Moving u from 0.05 to 0.30 — a year's work for most organizations — swings the steady state by 5.855. Moving from 0.70 to 1.00, which would require near-perfect orthogonality, swings it by 1.967. The first quarter of the range is worth three times the last. That is concavity, and it is a theorem.

Third — and this is the whole engineering argument — none of these rows required changing N, η, δ, or C. Same categories, same conversion, same leak, same cost. The entire column of steady states was produced by moving one parameter that costs nothing except the willingness to stop homogenizing.

9.6 Raising u — eight methods, with what each costs and what it buys

Ordered by speed of effect. Each entry gives the mechanism, the expected effect on u, the cost, the failure mode, and the measurement that confirms it worked.

M-1 . Silent-write-first Mechanism. Before any discussion of a question, every participant writes their independent answer and estimate, seals it, and submits it. Only then does discussion begin, and the sealed answers are read into the record first.

Effect on u. It does not raise the geometry; it stops the measured u from being destroyed by the cascade, which in practice is the same thing, since a collapsed discussion delivers the surplus of one member whatever the geometry. In bodies that have never done it, the gap between measured-under-cascade u and measured-blind u routinely exceeds 0.15 (Chapter 7, FM-7.5).

Cost. Three minutes per question. Zero capital. Zero headcount.

Failure mode. Performing it and then discarding it — writing the sealed answers and then proceeding as though the senior person's spoken opinion is the starting point. The sealed answers must be read into the record *before* discussion, in a randomized or reverse-seniority order, or the ritual is theatre.

Confirming measurement. Instrument 1 on the sealed answers and on the post-discussion positions; the gap is the cascade magnitude, tracked as a standing metric. This is the highest return-per-minute intervention in the manual, and adopting it requires no authority beyond the ability to say "everyone write it down first."

M-2 . Cut the shared incentive Mechanism. Identify the metric, target, or reward that every category is optimized against. Split it, stagger it, or replace it with category-specific objectives that are not derivable from one another.

Effect on u. Often the largest structural effect available. A common objective function is a **common mode**: it correlates the errors of otherwise independent instruments, exactly as a shared power supply correlates the noise of independent circuits. In Chapter 7's worked example, cross-lineage correlation ran at $\rho = 0.55$ among people from genuinely different traditions, the residual attributable to a shared funder with shared criteria; reducing it to $\rho = 0.40$ raised delivered surplus by 0.1550 nats with zero personnel change.

Cost. Political, sometimes severe. A single shared metric is administratively convenient and its removal is resisted by everyone whose job is administration.

Failure mode. Replacing one shared metric with another under a different name. Test it: if unit A's objective is a monotone function of unit B's, nothing changed.

Confirming measurement. Error correlation (Instrument 2) computed before and after, on scored outcomes. Expect the effect to take two to four periods to appear in the data, because the old incentive's influence on behaviour persists past its removal.

Related failure, named here because it is the same mechanism running the other way: Goodhart's Law and Campbell's Law. A measure that becomes a target ceases to be a measure. In this vocabulary: a shared target drives every category to the same point in state space, so u falls, so the measure loses the very diversity of view that made it informative. Goodhart capture is a u-destruction mechanism, and that is the sharpest available statement of why it happens.

M-3 . The adversarial role Mechanism. Assign, formally and rotationally, a member whose function for this decision is to build the strongest available case against the emerging consensus. The role is structural, not dispositional: it rotates, it is compensated, and its occupant is not held responsible for the position they argued.

Effect on u. Moderate to large, and immediate. It manufactures a non-parallel state where the geometry did not supply one — the cheapest way to obtain a nonzero u in a body that is genuinely homogeneous and cannot be reconstituted this quarter.

Cost. One person's time per decision, plus the meeting time to hear the case.

Failure mode. Two. **The role becomes a person** — one individual always plays it, acquires a reputation as difficult, and is discounted; rotation is what makes it structural rather than personal. Or **the role becomes a ritual** — the objection heard, thanked, minuted, ignored; require the case be answered in writing before the decision is recorded, and let an unanswered objection block it.

Confirming measurement. The fraction of decisions the adversarial case changed, even marginally. Zero over twenty decisions means the role is theatre.

Honest limitation. A manufactured u is not a measured u: the role produces an argument the body would not otherwise have heard, but the arguer's underlying state is still correlated with the body's. Badge it **[PARTIAL SUBSTITUTE]** and do not record it in a structural u estimate.

M-4 . The deliberate outsider Mechanism. Recruit one member from a training lineage with no overlap with any existing member, place them in a role with real decision rights, and protect

them from assimilation for a defined period.

Effect on u. Large, per Chapter 7's marginal table. The first outsider in a homogeneous body is worth most, because at low u the Lipschitz slope is steep: at N = 12, `df/du` at u = 0.1 is 4.74, so a 0.1 rise in u buys roughly 0.47 nats.

Cost. One hire, or one replacement. Real but bounded.

Failure mode — and it is the dominant one — assimilation. The outsider is hired, warmly welcomed, onboarded into the shared vocabulary, trained on the shared literature, evaluated against the shared metric, and within four quarters their error vector is indistinguishable from everyone else's. The organization concludes that hiring outsiders does not work, when what happened is that it converted the outsider into an insider and destroyed the asset it bought. §9.9 gives the protection protocol.

Confirming measurement. The outsider's error correlation against the body's mean, quarterly. A rising ρ is assimilation in progress, visible long before anyone notices it socially.

M-5 . Rotation Mechanism. Move people between categories on a fixed cadence — six to twenty-four months by the depth of the practice — with a required overlap period and a written handover.

Effect on u. Moderate, cumulative, durable. Rotation attacks u decay at its source: it stops an individual's state converging on their category's, and it manufactures the non-nested memberships Chapter 8 §8.10 identified as semilattice structure. Rotation is simultaneously a u intervention and a q intervention, which is why it appears in both chapters.

Cost. Real. Rotation costs productivity during transitions and depth permanently — a rotated practitioner is less expert in any one category than a fixed one. A genuine trade-off.

Failure mode. Rotating within a branch of the tree, which is nesting, not overlap, and buys nothing. Rotate **across** branches or do not rotate.

Confirming measurement. Pairwise u between categories, measured before rotation begins and after two full cycles. Also track the q statistic of Chapter 8 RN-8.

M-6 . Import a practitioner from an unrelated trade Mechanism. Bring in someone whose expertise is real, deep, and from a field with no institutional contact with yours: a logistician into a hospital, a stage manager into a research group, a bookbinder into a software archive project.

Effect on u. Potentially the largest single-person effect available, because lineage overlap is genuinely zero. From Table 7.2, an 80deg difference gives u = 0.9698 against 0.0670 at 15deg — a factor of fourteen.

Cost. Moderate to high in translation overhead: neither party can use the other's vocabulary. Chapter 20's translator role and boundary objects exist precisely for this.

Failure mode. Placing the outsider in an advisory position with no decision rights, where their non-parallel state has no channel into any decision. A state that cannot enter a decision contributes zero surplus to that decision regardless of its geometry. Advisory boards are the standard institutional method for acquiring u and simultaneously ensuring it is never spent.

Confirming measurement. Count decisions in which the imported practitioner's position appears in the decision record. Zero means the import is decorative.

M-7 . The red team Mechanism. A standing group, separately staffed and separately reporting, whose function is to attack the main body's conclusions — a permanent structure with its own membership, budget, and information access, rather than M-3's rotating role.

Effect on u. Large where it is real, which is rare. A red team reporting to the body it attacks is not a red team: the reporting line is a shared incentive (M-2) and will drive its error vector toward the main body's within a year.

Cost. High. It is a duplicate structure.

Failure mode. Capture by three routes — shared reporting line, shared recruitment pool, shared social environment. All three must be broken. A red team recruited from the main body's alumni is a mirror with a different letterhead.

Confirming measurement. Error correlation against the main body, quarterly. Below $\rho = 0.5$ the red team functions; above $\rho = 0.75$ it is captured and should be disbanded and rebuilt, because a captured red team is worse than none — it produces the appearance of adversarial review, which suppresses demand for the real thing.

M-8 . Recruit across training lineages, permanently Mechanism. Make lineage diversity a standing constraint on the pipeline, not a preference: cap the fraction of the body from any one training lineage, and require the lineage census (Chapter 7 Instrument 5, EC-9) run and published before any requisition is approved.

Effect on u. The largest **structural** effect available, and the slowest to arrive. Chapter 7's worked example computed the full ladder: three lineages at 5/4/3 gave u_bar = 0.5137; four lineages at 3/3/3/3 gave 0.5814; six lineages at 2x6 gave 0.6395; twelve distinct lineages gave 0.6975. Note the shape — moving from three lineages to twelve, the theoretical maximum, bought 0.2654 nats, while moving from one lineage to three would have bought far more. Concavity again. Get to three or four lineages and then stop optimizing lineage and start on incentives.

Cost. Recruitment friction. The pipeline is slower and candidates harder to evaluate, precisely because your evaluators are calibrated on your existing lineages — a homogeneous body's hiring process cannot correctly evaluate a candidate from outside its lineage, because evaluation is itself an instrument and the body has only one. Chapter 11 develops this.

Failure mode. Achieving lineage diversity and then homogenizing downstream through shared training, metrics, and review (FM-9.2). Chapter 7's FM-7.4 gives the diagnostic: a gap between the u the census predicts and the u error decorrelation measures means the force is downstream of hiring.

Confirming measurement. Both instruments, compared. Publish both.

9.7 Hong and Page: why this is not a fairness argument

Hong & Page (PNAS 2004) proved a result that is routinely misquoted and is worth stating precisely: under stated conditions, a randomly selected group of problem solvers can outperform a group composed of the best individual performers. The mechanism is that high-ability solvers

tend to share heuristics — they are good in similar ways — and therefore get stuck in the same local optima, while a more heuristically diverse group has more escape routes.

The result is a theorem under explicit assumptions about the problem space and the solver representations. It is not an empirical claim about all groups and it does not say ability is unimportant. It is a theorem with conditions, and the conditions matter.

What it supplies here is the answer to the objection raised against every intervention in §9.6: *"but the people we are excluding are better."* The answer is arithmetic. Take two candidates for a twelfth seat on a body at u_bar = 0.15, N = 12.

- Candidate A is measurably stronger individually and comes from the body's dominant lineage. Their error correlation with the body is ρ = 0.95, so their pairwise u against the body is $1 - 0.95^2 = 0.0975$.
- Candidate B is measurably weaker individually and comes from an absent lineage. Their error correlation with the body is ρ = 0.55, so their pairwise u is $1 - 0.55^2 = 0.6975$.

The body has 11 existing members, so the twelfth seat contributes 11 new pairs out of 66. Under the Roster Reduction:

```
with A:  u_bar = (55 x 0.15 + 11 x 0.0975)/66 = (8.250 + 1.073)/66 = 0.14126
          N_eff = 1 + 11 x 0.14126 = 2.554     S = 0.9377

with B:  u_bar = (55 x 0.15 + 11 x 0.6975)/66 = (8.250 + 7.673)/66 = 0.24126
          N_eff = 1 + 11 x 0.24126 = 3.654     S = 1.2960
```

Candidate B delivers 0.3583 more nats than candidate A — a 38% increase in the body's entire surplus input — and candidate A, the stronger individual, lowers the body's mean non-parallelism below its previous level. Hiring A makes the body worse in the only currency the law recognizes.

Say this plainly, because it is the sentence that will get the manual argued with: **the argument for non-parallelism is not about fairness, representation, or morale. It is about the arithmetic of a logarithm, and it would hold identically for sensors, models, or measuring instruments with no people involved.** That it also happens to support treating people decently is a convenience, not a premise; Chapter 22 takes up the ethics on its own terms.

9.8 How institutions destroy u without noticing

Almost every homogenizing mechanism below was installed deliberately, for a defensible reason, by someone competent. None was installed in order to destroy u. All of them do.

The unifying diagnosis: **u is not a property you acquire. It is a quantity that decays and must be maintained.** Unmaintained, it goes to zero at its own rate. That makes u a δ problem as much as a design problem, which is why Chapter 11 treats its maintenance as a first-class activity.

TABLE 9.3 — u DECAY. u_t = u_0 (1 − d_u)^t, starting from u_0 = 0.60

t	d_u=0.02	d_u=0.05	d_u=0.10	d_u=0.20
0	0.6000	0.6000	0.6000	0.6000
1	0.5880	0.5700	0.5400	0.4800
2	0.5762	0.5415	0.4860	0.3840

```
 3     0.5647     0.5144     0.4374     0.3072
 6     0.5315     0.4411     0.3189     0.1573
12     0.4708     0.3242     0.1695     0.0412
24     0.3695     0.1752     0.0479     0.0028
36     0.2899     0.0947     0.0135     0.0002
60     0.1785     0.0276     0.0011     0.0000
```

```
TABLE 9.4 — HALF-LIFE OF u UNDER CONVERGENCE RATE d_u
            t_half = ln 2 / ( -ln(1 - d_u) )

     d_u = 0.01     half-life = 68.97 periods
     d_u = 0.02     half-life = 34.31 periods
     d_u = 0.05     half-life = 13.51 periods
     d_u = 0.10     half-life =  6.58 periods
     d_u = 0.15     half-life =  4.27 periods
     d_u = 0.20     half-life =  3.11 periods
```

At a convergence rate of 10% per quarter — which is modest for a co-located team with a shared metric and a weekly all-hands — a body's non-parallelism halves in 6.58 quarters and is down to 5% of its initial value in five years. No one notices, because convergence feels like the team getting better at working together. It is, in fact, the team becoming one instrument.

Now the mechanisms.

D-1 . Hiring for culture fit The most efficient u-destroyer ever devised, and the one most likely to be described as a strength. "Culture fit" is operationally a screen for state similarity to the existing body: a correlation filter applied at entry, by evaluators who are themselves the reference vector.

Signature. Interview loops that include a "would I want to work with them" question. Candidates rejected for reasons that resolve, when written down, to unfamiliarity. High agreement among interviewers about candidates — which is itself a u = 0 measurement on the interview panel and should be read as such.

Correct replacement. Screen for **values fit** and **capability**, which are not similarity, and explicitly *against* lineage duplication; state in the requisition which lineages are over-represented. Where a "fit" judgement survives, require it in writing as a specific behavioural prediction — falsifiable — rather than as a feeling, which is a correlation reading.

D-2 . Promotion by similarity The same filter at every subsequent level, compounding. If each promotion round selects for similarity to the level above with even modest strength, after four levels the leadership is a near-pure sample from one lineage however diverse the entry cohort was. A diverse intake with a similarity-selecting promotion ladder produces a diverse workforce and a parallel leadership, which is the observed state of a great many institutions.

Signature. Lineage entropy falling monotonically with seniority. Compute the entropy of the lineage distribution at each level and plot against level; a downward slope is the diagnosis.

Correct replacement. Promotion criteria stated as observable outputs, not judgements of potential — potential is the channel through which similarity operates. Cross-level review panels drawn

from other branches. Publish lineage entropy by level annually.

D-3 . Shared training Everyone attends the same onboarding, the same courses, the same annual conference, and reads the same canonical texts. Done for excellent reasons — shared vocabulary lowers C substantially — and the cost is paid in u. A genuine trade-off, not a mistake; the engineering question is only where to put the boundary.

The rule this manual recommends: share the interface, not the model. Standardize the format in which claims, decisions, and artifacts are exchanged. Do not standardize the reasoning that produces them. A common file format costs nothing in u. A common analytical framework costs almost all of it.

Signature. Every member can complete every other member's sentence. New hires describe onboarding as "learning how we think here."

D-4 . Common incentives Covered as M-2 from the intervention side; from the destruction side, note how quietly it operates. A single organization-wide bonus metric, a single OKR cascade, a single set of promotion criteria — each is a common-mode input to every member's objective function. The members remain different people with different training and their **errors converge anyway**, because they are now all solving the same optimization problem.

Signature. The gap Chapter 7 FM-7.4 identifies: high lineage entropy, low measured u. If provenance says diverse and error decorrelation says parallel, look at the incentive first.

D-5 . Co-location without cross-assignment Putting people in one building raises contact rate. Contact without cross-assignment raises convergence without raising the overlap fraction q — a body that talks constantly and thinks identically, the maximum-convergence, minimum-structure configuration.

Signature. High interaction frequency, low measured u, and an org chart that is a strict tree (Chapter 8 FM-8.4). Open-plan offices with pure functional reporting are the canonical instance.

Correct replacement. Cross-assignment first, co-location second. Co-location amplifies whatever structure you have; if the structure is a tree, it accelerates the convergence. Chapter 12 gives the spatial version and Chapter 15 the threshold design that preserves separation where separation is load-bearing.

D-6 . Information cascades **Bikhchandani, Hirshleifer & Welch (1992)** showed that rational agents observing predecessors' actions can rationally abandon their own private information, at which point the sequence stops aggregating entirely — every later actor is copying, so no new information enters the record however many participate. In this vocabulary: **a cascade drives the observed u to zero regardless of the underlying geometry.** The body may hold twelve genuinely different instruments, and if the eleventh and twelfth speak after hearing the first ten, the reading is the first one's.

Signature. Strong dependence of stated position on speaking order. Rapid convergence in discussion. Unanimity arrived at quickly on questions that are genuinely hard.

Correct replacement. M-1, silent-write-first, an anti-cascade device by construction; Chapter 17's furniture and Chapter 22's protocols institutionalize it. A cascade is the cheapest u-destroyer to prevent and the most expensive to detect afterward, because it leaves a record of agreement.

D-7 . Preferential attachment and the homogenizing hub **Barabási & Albert (Science 1999)** showed that growth with preferential attachment produces scale-free degree distributions: a few nodes acquire very high degree. **Albert, Jeong & Barabási (Nature 2000)** showed such networks are robust to random failure and fragile to targeted attack on hubs.

The u consequence is a third property, specific to this manual's concern. A hub is connected to a large fraction of the network, and every node connected to it receives its input. Over time, the hub's state becomes a common-mode input to much of the body, and common-mode input correlates errors. The hub need not be coercive or even intentionally influential. It only needs to be listened to by many people.

Signature. One person, document, forum, or channel upstream of most positions. Ask twenty members where they formed their view on a contested question; if more than half trace to one source, that source is a homogenizing hub.

Correct replacement. Not removal — hubs carry real efficiency and robustness benefits. Install a second and third hub with genuinely different states, and ensure no member draws from only one. Two hubs at $\rho = 0.4$ with members drawing from both is a substantially different geometry from one, and it costs far less than dismantling the network. Note the interaction with Chapter 8: preferential attachment grows something close to a tree with a few high-degree centres, and §8.10 showed that tree structure drives a descriptor set to `N_eff = 1`. **Hub formation and tree formation are the same homogenization arriving by two routes.**

D-8 . Cohesion Cohesion is the most socially rewarded u-destroyer. A cohesive team is pleasant, low-friction, low-C, and it converges. Chapter 18 puts it as **cohesion is often the enemy of the collective intelligence factor c, because cohesion drives u toward 0**; the arithmetic version is that every mechanism making a group more comfortable with itself raises the convergence rate d_u unless specifically counteracted.

Signature. Rising satisfaction scores alongside falling disagreement rates. Meetings that end early. Departures described as "not a fit."

Correct replacement. Not to make the group unpleasant, which raises C without raising u, but to install counter-mechanisms — M-1, M-3, M-5 — that hold u against the cohesion gradient. Comfort and non-parallelism are not opposites, but neither are they automatically compatible; where they conflict, this manual chooses the arrangement that preserves both by design.

D-9 . The single review path Every output passes the same review, by the same reviewers, against the same rubric. A filter applied to all outputs makes all outputs more alike; over enough cycles the body learns to produce what passes, and what passes is one thing.

Signature. Convergence of output form over time. New members' work becoming indistinguishable from old members' work within a year. A rubric that has not changed in five years.

Correct replacement. Two independent review paths with genuinely different rubrics, or a panel whose composition rotates across lineages. And sunset the rubric: one unchanged in five years

is stored order consuming its own δQ maintenance while suppressing u — a case for Chapter 6's shed.

D-10 . Success The last and the hardest. A method that worked becomes the method. Everyone adopts it, because it worked. u falls, because it worked. **This is the u-side of Chapter 6's observation that success consumes its own headroom.** There, Q grew until δQ ate the budget; here, a successful approach propagates until it is the only approach. The two are the same event viewed from two parameters, and an institution in this condition has a rising Q and a falling u simultaneously, which is the fastest available route to h < 0.

Signature. One large success, then declining performance everyone attributes to changed external conditions. Universal endorsement of a specific method. Nobody able to describe what that method's failure conditions would look like.

Correct replacement. Reserve capacity for methods that are not the method — an explicit allocation, protected from the metric the successful method optimizes. It is expensive, and it is the intervention most likely to be cut in a downturn, which is exactly when it is most needed.

9.9 Protecting u after you have bought it

The most wasteful pattern in this subject: an organization spends a year and real money acquiring non-parallelism, then destroys it in four quarters through D-1 to D-10 without connecting the two events. Table 9.3 gives the arithmetic — from a hard-won u_0 = 0.60 at d_u = 0.10, the body is at 0.319 after six quarters and 0.170 after twelve. Translated to steady state:

TABLE 9.5 — REGRESSION AFTER THE CAMPAIGN STOPS
starting u = 0.57; N = 10, eta = 0.5, delta = 0.08, C = 0.30
Q_inf is the steady state the CURRENT u would sustain

quarters since	d_u = 0.05 u	Q_inf	d_u = 0.10 u	Q_inf	d_u = 0.15 u	Q_inf
0	0.5700	7.582	0.5700	7.582	0.5700	7.582
2	0.5144	7.050	0.4617	6.500	0.4118	5.931
4	0.4643	6.528	0.3740	5.461	0.2975	4.390
6	0.4190	6.016	0.3029	4.471	0.2150	2.979
8	0.3781	5.515	0.2454	3.536	0.1553	1.716
10	0.3413	5.025	0.1987	2.660	0.1122	0.613
12	0.3080	4.548	0.1610	1.848	0.0811	−0.325

At a convergence rate of 15% per quarter, an organization that reaches u = 0.57 and then stops maintaining it is back to an unsustainable steady state within three years. Not back to where it started — *below* where it started, because Q has grown in the meantime and δQ is now larger.

The protection protocol is the standing half of this chapter's Specification: EU-2, EU-6, EU-7, EU-9, and EU-10. Two points belong here rather than there.

Measure quarterly, never annually. The half-life at d_u = 0.10 is 6.58 quarters; an annual measurement samples a decaying signal at half its half-life and misses the trend for two years.

And when u falls, look for the mechanism, not the people. Every entry in §9.8 is structural. None requires anyone to behave badly and none is fixed by asking people to be more open-minded. **Exhortation is not an intervention on u; it is a request that people compensate personally for a machine that is running.** Chapter 20 makes the general form of this argument.

9.10 Measuring u — the standing protocol

Chapter 7 §7.7 gave the five instruments. The cadence that turns them into a standing measurement:

Quarterly. (1) *Blind item set* — twenty discriminating items from the body's actual work, fresh each quarter, simultaneous and sealed; discard unanimous items (EC-6), compute pairwise disagreement, convert to u, publish the matrix and the mean. (2) *Cascade gap* — re-administer five after discussion; blind u minus post-discussion u is the cascade magnitude, and its trend is the health of your elicitation. (3) *New-member correlation* per EU-7.

Annually. (4) *Error decorrelation on scored outcomes* (Instrument 2) — the ground truth, and the one that catches a body that argues about everything and is wrong in the same direction. (5) *Lineage census* (EC-9), entropy per provenance axis and per seniority level, which is D-2's diagnostic. (6) *Homogenizer inventory* per EU-9 — the single document that makes u destruction visible.

Reporting. Per EU-2, and always with an interval below u = 0.3 (EC-4): the Lipschitz slope makes point estimates misleading there.

The Specification

EU-1 — Silent-write-first on every consequential decision. Independent written positions, sealed, submitted before discussion and read into the record in randomized or reverse-seniority order. No exceptions for time pressure; the procedure costs three minutes.

EU-2 — Publish u quarterly with its interval. Format: `N = 10, u = 0.34 ± 0.06, N_eff = 4.06 [3.52, 4.60], S = 1.401, ceiling ln N = 2.303, deficit = 0.902`.

EU-3 — No shared objective function across categories. Each category's objective must be non-derivable from every other's. Test annually: if one category's objective is a monotone function of another's, they are one category for surplus purposes and must be recorded as such in N.

EU-4 — Rotate the adversarial role; never let it become a person. Rotation on a published schedule. The adversarial case must be answered in writing before the decision is recorded. An unanswered case blocks the decision.

EU-5 — Cap any single training lineage at one third of the body. Measured by the lineage census. A body exceeding the cap may not hire from the over-represented lineage until it is back in compliance. State the cap in the requisition.

EU-6 — Onboard to the interface, not to the model. Onboarding materials must be auditable against this distinction. Any module that teaches "how we think about X" is a homogenizer and goes in the register (EU-9) with an expiry date.

EU-7 — Protect new outsiders for eight quarters. Output-based evaluation, reviewed by someone outside the dominant lineage, with monthly error-correlation tracking. A ρ rising above 0.05 per quarter triggers a review of what is assimilating them. Onboard them to the interface, not the model, and do not evaluate them against insider criteria.

EU-8 — Give every member at least two upstream sources with measured ρ <= 0.5 between them. The anti-hub requirement (D-7). Audit by asking members to name where they formed their view on three contested questions.

EU-9 — Maintain the homogenizer register. Every shared rubric, metric, reading list, onboarding module, review path, and canonical forum, with install date, last re-justification, expiry, and the measured u across its population. No entry may exceed a three-year expiry unre-justified.

EU-10 — Budget u maintenance as an operating line. Not a project, not an initiative, not a campaign. A standing line covering the quarterly measurement, the rotation cadence, and the adversarial-role time.

EU-11 — Two independent review paths, or a rotating review panel. No output class may pass through a single fixed review path for more than three years.

EU-12 — Reserve capacity against the successful method. An explicit allocation, protected from the metric that method optimizes, for approaches that are not the current one. State it as a fraction and defend it in downturns specifically.

EU-13 — Badge manufactured u. Adversarial roles, red teams, and devil's-advocate protocols produce arguments, not geometry. Record them as `[PARTIAL SUBSTITUTE]` and do not enter them into a structural u estimate.

EU-14 — Test the binding constraint before spending on u. Compute `S / ln N` (Chapter 8, RN-6). Above 0.85, N binds and this chapter's interventions will not help. Below 0.85, u binds. Also compute h: **if h < 0, u interventions cannot rescue the system, because S is already capped at `ln N` and the deficit is in δQ or C.** Fix headroom first.

EU-15 — Never remove a person's ability to decline. Every intervention in §9.6 works by making a different position available, audible, and consequential. None may compel a person to hold a view, argue a position they reject, or remain in a role they wish to leave. The adversarial role is voluntary and compensated; rotation is opt-out-able with a stated reason. **A method that works by removing resistance has no coercive failure mode; its worst outcome is that someone declines.** Chapter 20 §20.6 develops this as a safety property of the whole approach.

The Failure Modes

FM-9.1 . THE MIRROR SECOND OPINION

Damages: u, and the credibility of every verification process built on it.

Clinical description. Verification is performed by a second party whose training, information, and incentives replicate the first's, and recorded as independent. Table 9.1 prices two reviewers at ρ = 0.95 at 13.4% of an independent pair, and at ρ = 1.00 at exactly zero.

Diagnostic signature. Second opinions agreeing above 95% of the time. Verification steps that have never once overturned an original finding. Reviewers drawn from the pool they review.

Detecting measurement. Error correlation between first and second party across scored outcomes. Anything above $\rho = 0.9$ means the verification is delivering under 26% of one independent check.

Fix. Draw the second party from a different lineage (M-8), give them different information access, and ensure they are not evaluated by the same person. If none is possible, record the verification as a consistency check, not an independent confirmation, and stop counting it as risk reduction.

FM-9.2 . ASSIMILATION OF THE OUTSIDER

Damages: u, on a delay of four to eight quarters after the acquisition.

Clinical description. The organization recruits genuinely non-parallel members, then converges their state onto the body's through onboarding, shared metrics, shared review, and social integration. The asset is destroyed by the machinery installed to make new members effective, and the organization concludes that hiring for difference does not work.

Diagnostic signature. New members' error correlation rising steadily through their first two years. Onboarding praised as "getting up to speed on how we do things." Outsiders who become the most enthusiastic exponents of the local view within a year.

Detecting measurement. EU-7's monthly tracking. Plot ρ against tenure for every member hired in the last three years; a rising curve is the diagnosis and its slope is the assimilation rate.

Fix. EU-6, EU-7, EU-9. The fix is not to onboard less; it is to onboard differently, and the distinction is entirely between teaching where things are and teaching how to think.

FM-9.3 . THE CAPTURED RED TEAM

Damages: u, plus the demand for real adversarial review.

Clinical description. An adversarial structure exists, is staffed and funded, and has converged on the body it was built to challenge — through shared reporting, shared recruitment, or shared social environment. It produces adversarial-shaped output confirming the body's position, which is worse than having no red team, because the structure suppresses demand for the function.

Diagnostic signature. Findings that have never changed a major decision; members recruited from the main body; reporting into the main body's leadership; correlation above $\rho = 0.75$.

Detecting measurement. Quarterly error correlation, plus the blunt count: how many decisions did red team review change?

Fix. Break all three capture routes at once — reporting line, recruitment pool, social environment — or disband. A partial fix leaves the appearance intact and the function absent.

FM-9.4 . EXHORTATION IN PLACE OF STRUCTURE

Damages: nothing directly, which is the problem. It consumes the appetite for intervention while changing no parameter.

Clinical description. A u problem is met with training on open-mindedness, workshops, statements of value, and requests that people speak up, while every mechanism in §9.8 keeps running. u does not move, and the failure is attributed to culture being hard to change.

Diagnostic signature. Interventions whose deliverable is a document or a session rather than a changed default, schedule, metric, or piece of furniture. Measured u unchanged four quarters after a major culture programme.

Detecting measurement. Measure u before and after; the mode is defined by producing no change, so the measurement is the detector.

Fix. Replace with a structural intervention from §9.6. **The test for whether an intervention is structural: could it work if nobody's attitude changed at all?** Silent-write-first passes. Cancelling the shared metric passes. A workshop does not. This is not an argument against psychological safety, which is a precondition for structure to function rather than a substitute for it: a body without it will not surface the sealed answers honestly, and a body with it and no structure will still cascade.

FM-9.5 . THE ADVISORY OUTSIDER

Damages: the return on a u acquisition, without damaging the measured u — which is what makes it hard to see.

Clinical description. A genuinely non-parallel member is placed where their state cannot enter a decision: an advisory board, a consulting arrangement, a research function with no operational authority, a report that is read and filed. The geometry is real and the channel is absent. Surplus requires interaction; a state that never interacts contributes zero.

Diagnostic signature. Advisory bodies whose recommendations appear in no decision record; consultants engaged repeatedly with no traceable change; research output described as "interesting."

Detecting measurement. Trace the last twenty decisions backward and count the fraction in which the outsider's position appears in the record at all, adopted or rejected. Below 20% is the diagnosis.

Fix. Decision rights, or at minimum a required-response protocol: the outsider's position must be answered in the decision record before the decision is final. This is M-3's requirement applied to M-6's asset.

FM-9.6 . CASCADE COLLAPSE

Damages: observed u, to zero, regardless of underlying geometry.

Clinical description. Per D-6. Positions form in sequence with visibility, so later positions carry no new information. The body may be geometrically rich and epistemically empty. Diagnostic signature: speaking-order dependence, rapid unanimity, cascade gap above 0.15. Detect with the cascade gap — blind u minus post-discussion u on the same items. Fix with EU-1; Chapter 11 gives the detection toolkit and Chapter 17 the furniture.

FM-9.7 . THE UNAUDITED HOMOGENIZER

Damages: u, continuously and invisibly.

Clinical description. A shared rubric, metric, reading list, onboarding module, or review path installed years ago for a good reason continues to run, unowned and unre-justified, converging every member it touches at a rate nobody measures.

Diagnostic signature. Ask who owns the rubric. If the answer is a role rather than a person, or the last revision predates half the current staff, the mode is present.

Detecting measurement. The homogenizer register (EU-9). The mode is defined by absence from the register, so building the register is the detection.

Fix. EU-9 with hard expiry dates: re-justify or retire. Retiring an obsolete rubric is a simultaneous u gain and Q shed — one of the few interventions that improves two terms of the headroom equation at once.

FM-9.8 . CHASING u AT NEGATIVE HEADROOM

Damages: the intervention budget, and the credibility of the framework.

Clinical description. A system with $h < 0$ invests heavily in u. Because $S \le \ln N$ and $h < 0$ means $\eta \ln N < \delta Q + C$, **no value of u whatsoever produces a nondecreasing step.** The intervention is well-designed, well-executed, correctly measured, and mathematically incapable of working — and the organization concludes the framework does not work.

Diagnostic signature. h computed and found negative, followed by a u programme. Or — far more common — a u programme launched without h ever being computed.

Detecting measurement. EU-14. Compute h first. Always.

Fix. Chapter 26's ordering: cut C, raise η, shed Q until $h > 0$. **Then** spend on u, and the same intervention that would have failed will work. The sequencing is not a preference; it is a feasibility condition from Proposition 4.1.

Worked Example

The engineering organization, eight quarters

A software engineering organization of roughly 90 people is arranged in ten functional categories: platform, application, data, infrastructure, security, quality, design, research, documentation, and operations. Its stated problem: "we ship, but we keep shipping the same class of mistake, and our post-incident reviews always conclude the same three things."

That last clause is an Ashby-shaped complaint (Chapter 8 §8.9), but the surprise log resolves it otherwise: the incidents are of familiar kinds, recurring. Not a missing category — a body producing one reading.

Step 1 — Survey Chapter 24's procedure, abbreviated to the parameters needed here.

Parameter	Value	Method
N	10	Declared functional categories. Relabeling audit (Chapter 8 §8.7): nine pass RN-4; “documentation” measures u = 0.11 against “quality” and is flagged but retained.
η	0.50	Chapter 4’s method. Design documents are written and kept, which is why η is not lower; no apprenticeship and no decision log, which is why it is not higher.
δ	0.08 per quarter	Chapter 3’s method, from decay across a two-quarter hiring freeze.
Q_0	3.0	Order proxy: runbook coverage, test suite health, documented-decision count, normalized.
C	0.30 per quarter	Chapter 5’s tally: cross-team coordination, on-call, release process, compliance.

Ceiling: η ln N = 0.50 x 2.3026 = 1.1513.

Step 2 — Measure u Blind item set: twenty architectural and prioritization judgment items, administered simultaneously and sealed. Six discarded as unanimous (EC-6); pairwise disagreement computed on the remaining fourteen and converted per Instrument 1, cross-checked against Instrument 2 on eighteen months of scored incident predictions.

```
    u_bar = 0.12  (interval [0.07, 0.17], reported as an interval per EC-4)
    N_eff = 1 + 9 x 0.12 = 2.08
    S     = ln 2.08 = 0.7324
    ceiling ln 10 = 2.3026
  redundancy deficit = 2.3026 - 0.7324 = 1.5702 nats  (68% of the ceiling forfeited)
```

Ten functional categories, ninety people, delivering the surplus of 2.08 categories.

Cascade gap: five items re-administered after discussion. Post-discussion u = 0.04 against blind u = 0.12. Gap of 0.08 — the discussion process is destroying two-thirds of what little non-parallelism the body has.

Step 3 — Check the binding constraint and the headroom (EU-14)

```
    S / ln N = 0.7324 / 2.3026 = 0.318      well below 0.85  ->  u binds, not N
```

```
h = eta ln N − delta Q − C
  = 1.1513 − 0.08 x 3.0 − 0.30
  = 1.1513 − 0.24 − 0.30
  = +0.6113                                POSITIVE
```

u binds and headroom is available — the configuration in which this chapter's interventions are applicable and feasible. Had h been negative, EU-14 and FM-9.8 would have routed the work to Chapter 26 instead.

Step 4 — The counterfactual What happens if nothing changes? Run the storage law forward at u = 0.12.

```
S = 0.7324, eta S = 0.3662, and delta Q + C at Q=3.0 is 0.24 + 0.30 = 0.54

t=0  Q = 3.000  ->  0.92 x 3.000 + 0.3662 − 0.30 = 2.826
t=1  Q = 2.826  ->  2.666
t=2  Q = 2.666  ->  2.519
t=3  Q = 2.519  ->  2.384
t=4  Q = 2.384  ->  2.259
t=6  Q = 2.145  ->  2.039
t=8  Q = 1.942  ->  1.853

Q_inf = (eta S − C)/delta = (0.3662 − 0.30)/0.08 = 0.827
```

The organization is converging to a steady state of 0.827 against a current 3.000 — a 72% reduction in stored order, over roughly 3τ = 37 quarters. Note that h is positive throughout: positive headroom means growth is *available*, not that it is *happening*. Corollary 3.2 gives the test — `Q_{t+1} > Q_t` iff `ηS_t > δQ_t + C_t`, and here `0.3662 < 0.54`. Declining.

Nobody would describe this as a crisis. It is a 4.5% quarterly erosion of runbook coverage, test health, and documented decisions, and it will be visible in about four years.

Step 5 — Design the campaign Three interventions, chosen by the §9.6 ordering and by what this body's diagnosis actually shows.

Q1 — M-1, silent-write-first, plus EU-1. The cascade gap of 0.08 says the discussion process is the largest immediate loss. Cost: three minutes per decision. Expected: recovery of most of the gap, and then more as people stop pre-converging in anticipation. `u: 0.12 -> 0.20` by Q2.

Q2 — M-2, cut the shared incentive. The audit finds one organization-wide quarterly delivery metric, bonus attached, against which all ten categories are assessed — security, quality, and research all optimizing a throughput number. Replace with three clusters carrying non-derivable objectives (EU-3): delivery, reliability, discovery. `u: 0.20 -> 0.35` over three to four quarters.

Q3 — M-3 and M-8 together. A rotating adversarial role on every architectural decision (badged `[PARTIAL SUBSTITUTE]` per EU-13, so not entered into the structural estimate), plus two requisitions under a lineage constraint: the census shows 61% of engineering staff trained in one framework tradition, exceeding EU-5's one-third cap. Expected: `u: 0.35 -> 0.50` over four quarters.

Not attempted: raising N. `S / ln N = 0.318` says N does not bind. The flagged "documentation" category, measuring u = 0.11 against "quality," is left in place and re-tested at Q8.

Step 6 — The trajectory Run the storage law forward with the campaign's u path. Parameters unchanged: `η = 0.5`, `δ = 0.08`, `N = 10`, `C = 0.30`, `Q_0 = 3.0`.

TABLE 9.6 — CAMPAIGN TRAJECTORY, SIXTEEN QUARTERS

t	u	N_eff	S = ln N_eff	h_t	Q_t	->	Q_{t+1}
0	0.12	2.08	0.7324	+0.6113	3.000	->	2.826
1	0.12	2.08	0.7324	+0.6252	2.826	->	2.666
2	0.20	2.80	1.0296	+0.6380	2.666	->	2.668
3	0.28	3.52	1.2585	+0.6379	2.668	->	2.784
4	0.35	4.15	1.4231	+0.6286	2.784	->	2.972
5	0.41	4.69	1.5454	+0.6135	2.972	->	3.207
6	0.46	5.14	1.6371	+0.5947	3.207	->	3.469
7	0.50	5.50	1.7047	+0.5737	3.469	->	3.744
8	0.52	5.68	1.7370	+0.5518	3.744	->	4.013
9	0.54	5.86	1.7681	+0.5302	4.013	->	4.276
10	0.55	5.95	1.7834	+0.5092	4.276	->	4.526
11	0.56	6.04	1.7984	+0.4892	4.526	->	4.763
12	0.56	6.04	1.7984	+0.4703	4.763	->	4.981
13	0.57	6.13	1.8132	+0.4528	4.981	->	5.189
14	0.57	6.13	1.8132	+0.4362	5.189	->	5.381
15	0.57	6.13	1.8132	+0.4208	5.381	->	5.557
16	0.57	6.13	1.8132	+0.4068	5.557	->	5.719

```
Q_inf at u* = 0.57:  (0.5 x 1.8132 - 0.30)/0.08 = 7.582
tau = 1/delta = 12.5 quarters;  3tau = 38,  5tau = 62
```

Four things to read from this table.

First, the turnaround is at t = 2, not t = 0. Q continues to fall for two quarters after the campaign begins, because the integrator is still discharging its initial transient `(1−δ)^t Q_0` while the driven term builds. This is Book One's Node 09 RECURSE and Chapter 1's central point: **a leaky integrator does not respond instantly, and a campaign evaluated at one quarter will be judged a failure at exactly the moment it is working.** Anyone who cancels this programme at Q2 because "the numbers are still going down" has misread the circuit.

Second, the arrival is slow. τ = 12.5 quarters, so 3τ ~ 38 quarters to reach 95% of the new steady state. At t = 16 the system is at 5.719 against a fixed point of 7.582, 75% of the way. The intervention was fast; the arrival is slow, and the slowness is δ, not the intervention.

Third, headroom falls monotonically as Q rises — from +0.6113 at t = 0 to +0.4068 at t = 16. Chapter 6's observation that success consumes its own headroom, arriving on schedule. At the fixed point, headroom is `1.1513 − 0.08 x 7.582 − 0.30 = 0.2447`: positive, but 60% consumed. The maintenance transition lands at about Q = 10.6, where `δQ + C = η ln N` exactly, and the organization should know that now rather than discovering it in a decade.

Fourth, compare against the counterfactual. At t = 8, campaign Q = 4.013 against counterfactual 1.853, a factor of 2.17; at the fixed points, 7.582 against 0.827, a factor of 9.17. The campaign cost three minutes per meeting, one metric redesign, one rotating role, and two lineage-constrained hires — and it moved the organization's stable stored order by a factor of nine.

Step 7 — What was and was not achieved Honest accounting. What did **not** change: N stayed at 10, η at 0.50, δ at 0.08, C at 0.30, and headcount rose by two on a base of ninety. Every gain in Table 9.6 came from one parameter, moved mostly by cancelling things.

What was not delivered: the final u of 0.57 is well short of 1.0, and `S = 1.8132` is 79% of the ceiling. By RN-6's test `S / ln N = 0.787`, still below 0.85, so u binds marginally and a little remains. Take it carefully — from Table 9.2, u from 0.60 to 0.70 buys 0.0658 nats of ηS and 0.82 of `Q_inf`, against 0.1939 and 2.42 for the move from 0.12 to 0.20. The campaign has collected the front-loaded portion, and the correct next move is η, not more u.

Step 8 — Maintenance, or it all comes back The campaign's output is `u* = 0.57`. Per §9.9, that figure decays. The organization must now budget maintenance as a standing line (EU-10) or repeat Table 9.5:

```
at d_u = 0.05:  Q_inf sustainable falls  7.582 -> 4.548 over 12 quarters
at d_u = 0.10:  Q_inf sustainable falls  7.582 -> 1.848 over 12 quarters
at d_u = 0.15:  Q_inf sustainable falls  7.582 -> -0.325 over 12 quarters
```

At a 15% quarterly convergence rate the organization is back below its unsustainable threshold in three years, carrying a Q of around 6 when it gets there — so the second decline is steeper than the first. Standing commitments per EU-2 and EU-7 to EU-11: quarterly measurement, new-member correlation tracking, the homogenizer register with three-year expiries, the rotation cadence, and two independent review paths.

Step 9 — Falsification (Chapter 25) Stated in advance, before the campaign begins, and frozen.

- **The blind u must rise.** If measured u at Q4 is not above 0.25, the incentive change did not affect the geometry and the model of what was homogenizing this body is wrong — most likely the shared metric was not the binding common mode and shared hiring is, given the 61% concentration.
- **The trajectory must turn by Q3.** Table 9.6 puts the minimum at t = 2, `Q ~ 2.666`. Still falling at t = 4 means δ was mis-estimated or the proxy is not tracking Q.
- **Q at t = 8 must be within measurement error of 4.013.** A miss on the level with a correct turning point implicates η; a miss on both implicates δ.
- **The order proxy must not be a repackaging of the input.** Paper E §6.3. If documented-decision count is partly a measure of how much the body interacts, it is contaminated and must be dropped from the composite. **This is the most common measurement error in the framework**; Chapter 24 §24.7 gives the test.
- **The shuffled-input control is mandatory.** Refit `Y_{t+1} = aY_t + bF(U_t) − Ĉ_t + ε_t` with the U series randomly permuted in time. If the shuffled fit performs comparably on held-out data, the apparent link between this organization's geometry and its stored order is an artifact and **this entire analysis must be withdrawn, including the parts that appeared to work.** It is the control most often omitted and the one that would kill this result.

u is the fast parameter. It collects on every category already bought, it moves in a quarter, and most of the moves are subtractions — cancel the shared metric, the sequential discussion, the culture-fit screen; sunset the five-year-old rubric. It is also the parameter that decays fastest and most invisibly, because every force acting on it — cohesion, success, shared training, a good hire being welcomed — feels like the organization working well.

Non-parallelism is not a state an institution achieves. It is a load an institution carries, and the moment it stops carrying it the load slides back down. Chapter 10 splits the load into its two independent components, so you can tell which half you are carrying and which half you dropped.

Chapter 10 . Crossing and Divergence

Chapter 9 treated u as one quantity and asked how to raise it. Useful, but not sufficient for design. An engineer who knows only a group's total non-parallelism knows its output and not its supply chain. Two organizations can present the same u, the same `N_eff`, the same surplus, and be built of completely different materials — and when one fails, it fails in a way the other cannot.

This chapter splits u into its two structural sources and shows how to buy, measure, and maintain each separately. It explains why an institution can hire across every discipline on the register and still produce decisions indistinguishable from those of a single person.

The Law

10.0.1 The theorem, stated exactly

Paper A, Theorem 4.1 (two-source decomposition). For a pure block vector,

```
u = u_cross + u_div

    u_cross = beta^2              (between-category component)
    u_div   = alpha^2 sin^2(gamma) (within-category component)

max( f(u_cross), f(u_div) )  <=  f(u)  <=  f(u_cross) + f(u_div)
```

with `f(u) = ln(1 + (N−1)u)` as established by Paper A Theorem 2.1 under axioms S1-S4, and `alpha^2 + beta^2 = 1`.

That is the whole formal content of this chapter. Everything below is either a reading of those five lines or an instrument for measuring their terms.

10.0.2 What alpha, beta, and gamma are

Take two participants, x and y, each a unit state vector in a space carrying a block structure — the blocks being the N categories the survey (Chapter 24) declared. Let x lie inside one block. Decompose y relative to that block:

```
y = alpha * (unit vector inside x's block, at angle gamma from x)
  + beta  * (unit vector outside x's block)

alpha^2 + beta^2 = 1,  alpha, beta >= 0
```

Then the inner product only survives on the shared block:

```
x . y   = alpha cos(gamma)
u       = 1 − (x . y)^2
        = 1 − alpha^2 cos^2(gamma)
        = 1 − alpha^2 + alpha^2 sin^2(gamma)
        = beta^2 + alpha^2 sin^2(gamma)
        = u_cross + u_div
```

A one-line verification of the theorem's form, not a re-derivation — worth doing once, because it fixes what the three parameters mean in an organization:

Symbol	Formal meaning	Organizational reading
`beta^2`	squared component of y outside x's category block	the fraction of the second party's competence that the first party cannot evaluate
`alpha^2`	squared component of y inside x's block, `= 1 − beta^2`	the shared ground: the part of the work both parties can actually judge
`gamma`	angle between the two inside the shared block	how much they disagree about the things they both understand

`u_cross` is bought by hiring across; `u_div` is bought by permitting disagreement within. Different transactions, different costs, different failure modes — and, as §10.3 shows, a hard budget relation between them that almost nobody accounts for.

10.0.3 The two bounds, and what each one licenses

The bound is two-sided and the two sides do different jobs.

Lower bound — the guarantee.

```
f(u) >= max( f(u_cross), f(u_div) )
```

Whatever you already had from the stronger source, adding the other cannot take it away. This is the safe half of the theorem, and it licenses one design statement: *adding a genuinely different source is never negative.*

Upper bound — the ceiling, and the warning.

```
f(u) <= f(u_cross) + f(u_div)
```

This follows from subadditivity on u (Paper A, Theorem 3.4), itself a consequence of strict concavity (3.3). The combined system is worth **at most** what running the two sources separately and adding their surpluses would give.

Read that again in its practical form. **The theorem promises no synergy beyond additivity, and honest design must not assume any.** No clause in Paper A says a cross-disciplinary team with internal dissent produces more than the sum of a cross-disciplinary team and a dissenting team. The mathematics forbids more than the sum. It does not forbid less, and — because `f` is strictly concave — the truth is systematically *less* whenever both sources are nonzero. How much less is computable, and it is not a rounding error.

```
N = 8,  f(u) = ln(1 + 7u)

split (u_cross : u_div)   u      f(u) true   upper bound   overstatement
-----------------------------------------------------------------------
50 / 50                  0.20     0.8755       1.0613          21.2%
```

50 / 50	0.40	1.3350	1.7509	31.2%
50 / 50	0.60	1.6487	2.2628	37.3%
50 / 50	0.80	1.8871	2.6700	41.5%
50 / 50	1.00	2.0794	3.0082	44.7%
75 / 25	0.20	0.8755	1.0179	16.3%
75 / 25	0.60	1.6487	2.1409	29.9%
75 / 25	1.00	2.0794	2.8442	36.8%

At a balanced split and full non-parallelism, the upper bound overstates real surplus by 44.7 percent. **Any plan that adds the two sources' surpluses overstates its deliverable by up to forty-five percent, and the error grows exactly where the ambition is highest.** The upper bound is a proof device, not a plan.

10.0.4 The marginal identity

Because `u = u_cross + u_div` is a plain sum, and `f` depends on u alone (axiom S1, geometric dependence),

```
d f / d u_cross  =  d f / d u_div  =  (N − 1) / (1 + (N−1) u)
```

At the margin, in the parameter u, a unit of crossing and a unit of divergence are worth exactly the same. There is no premium on cross-disciplinarity in this law. The theorem is indifferent to where the non-parallelism came from.

This matters more than it looks. Any organizational preference for cross-hiring over internal dissent is a preference about *acquisition cost and observability*, not about value: a procurement decision dressed as an epistemic one. Chapter 26's lever ranking must be built on that distinction and not on the fiction that crossing is intrinsically superior.

10.0.5 Badges

- **Proven (Paper A, Thm 4.1):** the decomposition and both bounds, for a pure block vector.
- **Proven (Paper A, Thm 3.3, 3.4):** strict concavity and subadditivity on `u`, from which the upper bound follows.
- **Design axiom, not derived (S3):** the affine effective count `g(u) = e^{f(u)} = 1 + (N−1)u`. Everything here that speaks of `N_eff` inherits that choice, made because it renders the effective count additive in difference — what an engineer wants from a headcount audit. Say so whenever you lean on it, as this book does here.
- **Postulate, badged:** that participants can be represented by state vectors in a block-structured space at all, and that the declared categories correspond to the blocks. A modeling assumption, not a result, falsifiable in the manner of Chapter 25: shuffle the category assignment, and if the fitted parameters do not degrade, the block structure was not real.
- **Honesty clause, restated:** the triangle inequality FAILS for geometric triples (Paper A §5). Subadditivity holds on `u`, and only on `u`. Do not chain three participants and expect a metric.
- **Honesty clause, restated:** as deployed, the surplus functional is a **lower bound** on the ensemble's information, not the information itself. We log the floor, not the ceiling. A group with `u_div` near zero may still hold information this instrument cannot see; it will not hold

information the instrument *credits*, and stored order (Chapter 23) responds to what the instrument credits.

- **Honesty clause, restated:** Paper A is pure mathematics and makes no empirical or metaphysical claim. Nothing here says a cross-disciplinary team is *good*. It says what a specified functional returns.
- **Refused (CANON §5.4):** no protocol in this chapter removes a person's ability to decline to disagree. Divergence performed on demand is not divergence but a second kind of compliance, and it measures as noise.

10.0.6 The structural constraint

One consequence of `alpha^2 + beta^2 = 1` appears nowhere in the theorem's headline and governs almost every real design:

```
u_div = alpha^2 sin^2(gamma) <= alpha^2 = 1 − beta^2 = 1 − u_cross

so   u_cross + u_div <= 1   (as it must, u being in [0,1])
and  max achievable u_div, given u_cross  =  1 − u_cross
```

The two sources are not independent supplies from separate accounts; they compete for the same unit budget. At `u_cross = 0.90`, no amount of internal argument can buy more than `u_div = 0.10` for that pair, even at total within-block disagreement (`gamma = 90 degrees`). The ground on which two people can disagree is exactly the ground they hold in common, and crossing spends it.

This is the mathematical form of an observation every practitioner has made: the more foreign two specialists are to each other, the less able they are to have a real argument. They talk past each other instead — `beta^2` without `gamma`, which the law credits correctly as non-parallelism but which is a different product from a disciplined disagreement and cannot substitute for one when the question is internal to a discipline.

The Engineering

10.1 Two ways not to be parallel

Case A — the crossing. A structural engineer and a public-health epidemiologist assess a proposed bridge. The epidemiologist cannot evaluate the fatigue calculations; the engineer cannot evaluate the ambulance-transit consequences of a lane closure. Almost all of the second party's state lies outside the first's block: `beta^2 = 0.85`. Within the sliver they share — cost, schedule, public communication — they broadly agree, `gamma = 12 degrees`.

```
u_cross = 0.85
u_div   = (1 − 0.85) ∗ sin^2(12 deg) = 0.15 ∗ 0.0432 = 0.0065
u       = 0.8565
```

Case B — the divergence. Two structural engineers, same degree programme, same code books, assess the same bridge. `beta^2 = 0.03`. But one is a fatigue specialist who believes the connection detail will crack in twenty years and the other does not. Within the shared block they are 65 degrees apart.

```
u_cross = 0.03
u_div   = 0.97 * sin^2(65 deg) = 0.97 * 0.8214 = 0.7968
u       = 0.8268
```

At `N = 6`, `f(u) = ln(1 + 5u)`: Case A gives `ln(5.2825) = 1.6644`, Case B gives `ln(5.1340) = 1.6360` — within two percent of each other. **The law prices a foreign expert and a determined internal dissenter as near-equivalent goods, and an organization that buys only one of them has bought half a portfolio at full price.**

The difference is not in the surplus but in what the surplus is *about*. Case A's non-parallelism lives entirely in the question "what else does this bridge touch?" Case B's lives entirely in "is this bridge sound?" An organization with only Case A never learns that the connection detail cracks. An organization with only Case B builds a structurally impeccable bridge that closes the only route to the hospital.

10.2 Reading beta and gamma as organizational facts

Each parameter has a direct organizational correlate, and each is acquired by a different administrative act.

	`u_cross` (beta^2)	`u_div` (alpha^2 sin^2 gamma)
acquired by	recruiting, seconding, rotation across disciplines	protecting dissent, rotating adversarial roles, independent estimates first
visible in	org chart, CV, hiring record	minutes, dissent log, the error-decorrelation audit
costs	salary, onboarding, translation burden (C)	social friction, slower meetings, managerial nerve
decays via	attrition, assimilation to the dominant discipline	cascades, promotion by similarity, shared incentives
audit interval	annual (slow-moving)	quarterly at most (fast-moving; see Chapter 11)

The asymmetry in the "visible in" row is the engine of the pathology this chapter exists to diagnose. `u_cross` is legible on a spreadsheet: it survives an audit and can be reported to a board. `u_div` exists only in behaviour and leaves a record only if someone deliberately builds one. **An institution that manages what it can see will accumulate `beta^2` and let `gamma` go to zero without ever making a decision to do so.**

10.3 The budget constraint nobody accounts for

The ceiling `u_div <= 1 - u_cross` changes what "more diversity" means at different operating points, and its shape is counter-intuitive enough to need a table. Estimate `sin^2(gamma)` from the correlation `r` of two parties' independent judgments on items *both* can evaluate: `sin^2(gamma) = 1 - r^2` (an estimator choice, badged; see §10.10). Then `u_div = (1 - u_cross)(1 - r^2)`.

```
u_cross |  r=0.95        r=0.90        r=0.60        r=0.00
```

```
        |  u_div  u      u_div  u      u_div  u      u_div  u
--------+-------------------------------------------------------
 0.00   |  .098  .098    .190  .190    .640  .640   1.000 1.000
 0.10   |  .088  .188    .171  .271    .576  .676    .900 1.000
 0.25   |  .073  .323    .142  .392    .480  .730    .750 1.000
 0.50   |  .049  .549    .095  .595    .320  .820    .500 1.000
 0.72   |  .027  .747    .053  .773    .179  .899    .280 1.000
 0.90   |  .010  .910    .019  .919    .064  .964    .100 1.000
```

Read the rows. At `u_cross = 0.90`, driving within-block correlation from 0.95 to 0.00 — near-total agreement to total independence — moves u from 0.910 to 1.000. At `u_cross = 0.00`, the identical change in `r` moves u from 0.098 to 1.000. At `N = 6`:

```
u_cross = 0.90:  f(1.000) − f(0.910) = 1.7918 − 1.7138 = 0.0780
u_cross = 0.00:  f(1.000) − f(0.098) = 1.7918 − 0.3988 = 1.3930
```

The same dissent intervention is worth nearly eighteen times more in a monodisciplinary group than in a heavily cross-staffed one — not because dissent is less valuable there in principle, but because the cross-staffing has already consumed most of the unit budget, leaving little shared ground on which the dissent can register. The converse holds too, and is the correction to the usual enthusiasm.

10.4 Which lever is live — marginal productivity in the design variables

In the *parameter* u, both sources are worth the same at the margin (§10.0.4). But the design variables an engineer controls are `beta^2` (who you hire) and `sin^2(gamma)` (what you permit), and in those the marginal productivities are not equal:

```
u = beta^2 + (1 − beta^2) sin^2(gamma)

d u / d(beta^2)       = 1 − sin^2(gamma) = cos^2(gamma)
d u / d(sin^2 gamma)  = 1 − beta^2       = alpha^2
```

```
beta^2  sin^2 g |  du/d(beta^2)   du/d(sin^2 g)  |   u
----------------+-------------------------------+--------
 0.00    0.00   |     1.000           1.000       |  0.000
 0.00    0.50   |     0.500           1.000       |  0.500
 0.00    0.90   |     0.100           1.000       |  0.900
                |                                 |
 0.50    0.00   |     1.000           0.500       |  0.500
 0.50    0.50   |     0.500           0.500       |  0.750
 0.50    0.90   |     0.100           0.500       |  0.950
                |                                 |
 0.72    0.00   |     1.000           0.280       |  0.720
 0.72    0.50   |     0.500           0.280       |  0.860
 0.72    0.90   |     0.100           0.280       |  0.972
                |                                 |
 0.90    0.00   |     1.000           0.100       |  0.900
 0.90    0.50   |     0.500           0.100       |  0.950
 0.90    0.90   |     0.100           0.100       |  0.990
```

The rule that falls out is symmetric, and it is the single most useful operational statement in this chapter:

```
RULE 10.4  (which lever is live)

  Hiring across is most productive when the group already agrees internally:
    du/d(beta^2) = cos^2(gamma)  -> maximal at gamma = 0.
  Permitting dissent is most productive in a monodisciplinary group:
    du/d(sin^2 gamma) = alpha^2  -> maximal at beta^2 = 0.

  Therefore the correct first move is always the one you have LESS of.
```

Each lever's productivity is highest exactly where the other has not been pulled. A self-correcting rule: you cannot go far wrong buying the scarcer source, and you will go wrong buying more of the one you already have.

An organization at `beta^2 = 0.72, sin^2(gamma) = 0.00` — the classic case this chapter diagnoses — gets `du/d(beta^2) = 1.000` from another cross-hire against `du/d(sin^2 gamma) = 0.280` from permitting dissent. The naive reading is: keep cross-hiring, it is nearly four times as productive per unit. That reading is wrong twice over.

First, the units are not comparable. Moving `beta^2` by 0.10 requires restructuring a team; moving `sin^2(gamma)` by 0.28 requires a meeting protocol (Chapter 17) and the nerve to keep it. Chapter 26's cost ranking converts marginal productivity into value per unit of expenditure, and by that measure the protocol wins in almost every real case.

Second, surplus is not the whole brief (§10.1). The two components are surplus about *different questions*, and a group at `beta^2 = 0.72, gamma = 0` has bought a wide field of view with no resolution anywhere in it. Chapter 7's phrase applies with force: **three honest windows set wide capture most of what a gallery could — but a window that reports what the window next to it reports is not a second window.**

10.5 What the upper bound actually licenses

The upper bound is stated so often as "subadditivity" that its practical content is easy to lose. Three statements, each forbidding a specific real mistake.

Statement 1 — no synergy is promised. The theorem gives a ceiling and no floor above the max. Any memo arguing that cross-disciplinary staffing *and* structured dissent together "will be greater than the sum of its parts" makes a claim the mathematics does not support and in fact bounds against. **Honest design assumes additivity at best and plans against the true concave value.**

Statement 2 — the true value is systematically below the sum. Strict concavity (Thm 3.3) makes the inequality strict whenever both components are nonzero. The overstatement table in §10.0.3 quantifies it. Use `f(u_cross + u_div)`. Never use `f(u_cross) + f(u_div)`.

Statement 3 — the bound is not tight, so it is not an estimate. A loose upper bound is a proof of impossibility, not a forecast: it says where surplus *cannot* be, and nothing about where it is.

A fourth statement, badged as this book's interpretation rather than a result: **because the deployed functional is a lower bound on the ensemble's information, a group may realize benefits from combining the two sources that this instrument does not credit.** A reason to build

such groups; not a reason to write the uncredited benefit into a plan. We log the floor and spend against the floor.

10.6 The lower bound is the floor you can actually spend

```
f(u) >= max( f(u_cross), f(u_div) )
```

The lower bound has a use the upper bound does not: it is a *guarantee*, and guarantees are what you commit to. Its tightness varies, and knowing where it is tight tells you when the second source is worth the trouble. The gain of the true value over the best single source, at `N = 8`:

```
u     u_cross  u_div   max f (floor)   f(u) true   gain over best single
-------------------------------------------------------------------------
1.00   0.500   0.500      1.5041        2.0794     +0.5754   (+38.3%)
1.00   0.700   0.300      1.7750        2.0794     +0.3045   (+17.2%)
1.00   0.900   0.100      1.9879        2.0794     +0.0916   (+4.6%)
1.00   0.980   0.020      2.0618        2.0794     +0.0177   (+0.9%)
1.00   1.000   0.000      2.0794        2.0794     +0.0000   (0.0%)
```

The bound is tight exactly when the portfolio is lopsided, and slack exactly when it is balanced — which makes the lower bound a direct readout of how unbalanced your two sources are. At a 50/50 split the true value exceeds the guaranteed floor by 38 percent; at 98/2 it exceeds it by under one percent, which is to say the minority source has bought nothing you did not have.

That gives a diagnostic requiring no new instrumentation:

```
BALANCE INDEX

  B := [ f(u) − max( f(u_cross), f(u_div) ) ] / f(u)

  B near 0      one source does all the work; the other is decoration
  B in 0.1–0.3  a real but lopsided portfolio
  B > 0.3       both sources materially loaded
```

For the rows above: B = 0.2767 at 50/50, 0.1464 at 70/30, 0.0440 at 90/10, 0.0085 at 98/2. Compute B every audit cycle — it needs no data beyond the decomposition you already have, and it tells a board whether its diversity programme has one leg or two.

10.7 The classic pathology: rich in crossing, starved of divergence

This is the central practical diagnosis of the chapter, and it deserves to be stated without hedging. A very large number of contemporary organizations — institutes, consultancies, agencies, foundations, product teams, review panels, standards bodies — have executed a genuine and expensive programme to raise `u_cross` and simultaneously, usually without noticing, driven `u_div` toward zero. The mechanism is not conspiratorial. It is administrative, and every step is locally reasonable.

```
STEP                                          ON u_cross   ON u_div
-------------------------------------------------------------------
1. Recruit deliberately across disciplines   UP           neutral
```

```
2. Onboard everyone through one induction    neutral      DOWN
3. Adopt one shared framework/vocabulary     neutral      DOWN
4. Tie all bonuses to one composite metric   neutral      DOWN (hard)
5. Promote the people who integrate well     neutral      DOWN
6. Report diversity as headcount by          reported UP  invisible
   category; never report disagreement
7. Escalate disputes to one decision-maker   neutral      DOWN (hard)
-----------------------------------------------------------------
```

Steps 2 through 5 are the things you do to make a cross-disciplinary team *function*: its `C` term (Chapter 5) is enormous because translation is expensive, and shared vocabulary, shared metrics and clear escalation are the cheapest ways to cut it. **The organization pays for its `u_cross` by spending its `u_div`, and the transaction never appears on any ledger because only one of the two quantities is measured.**

The signature is unmistakable once you look for it. The org chart shows six disciplines and the last forty decision memos show no recorded dissent. Everyone can explain why their discipline matters, and no two people in the same discipline have ever publicly disagreed about a substantive call. Reviews described as "robust" produce, on inspection, a 100 percent approval rate with comments. "We're all aligned on this" is a report of success. New arrivals from outside the dominant discipline stop objecting within four to eight months, and this is described as *settling in*.

Run the arithmetic at `N = 8`, against three alternative portfolios:

```
N = 8,  f(u) = ln(1 + 7u)

portfolio              u_cross  u_div   u     f(u)    N_eff   floor   ceiling
-------------------------------------------------------------------------
----------
diverse hires, one opinion  0.750  0.000  0.750  1.8326   6.250  1.8326  1.8326
  + a dissent protocol    0.750  0.050  0.800  1.8871   6.600  1.8326  2.1327
monodiscipline, real fight  0.000  0.750  0.750  1.8326   6.250  1.8326  1.8326
balanced                 0.375  0.375  0.750  1.8326   6.250  1.2879  2.5757
```

Three of these four rows have identical surplus. That is the law being honest: it counts non-parallelism and does not care about its provenance. But look at the floor and ceiling columns. The first row is *pinned* — floor equals true value equals ceiling, because one component is zero. No slack, no second source, no redundancy in its difference: a single-channel instrument with six labels on it. Remove its one source and it holds nothing. The balanced row has the same output today and, if either source were removed, would still hold 1.2879.

A portfolio with the same surplus and two live sources is strictly more robust than a portfolio with the same surplus and one, and the surplus number alone cannot tell them apart. This is the Albert-Jeong-Barabási result in another domain: a structure concentrated on one hub is robust to random failure and fragile to targeted attack on that hub. A monoculture of opinion, however cross-disciplinary its headcount, has exactly one hub.

10.8 The mirror pathology: the monoculture that argues

The opposite failure is rare in large institutions and endemic in small expert ones: u_div high, u_cross near zero. Fourteen economists who genuinely and vigorously disagree. A surgical team that argues every case. Twelve cryptographers who are each other's fiercest critics. gamma large, beta^2 = 0.02.

u_cross = 0.02, u_div = 0.98 * sin^2(70 deg) = 0.98 * 0.8830 = 0.8654
u = 0.8854, N=8: f(u) = ln(7.1978) = 1.9738, N_eff = 7.198

Excellent surplus — about a strictly bounded set of questions, because no component of anyone's state lies outside the block. Every question the block does not contain returns u = 0: not low surplus, *no* surplus, because there is no non-parallel component to register it.

The signature: fierce, high-quality internal debate and a reputation for rigour; decisions impeccable within the frame and catastrophic outside it; failures always described afterwards as "outside our remit"; recruiting that selects for people who can survive the internal debate, which is a filter for people already inside the block. Ashby's Law of Requisite Variety applied and failed — variety deep in one dimension and zero in all others absorbs only the disturbances that arrive along that dimension.

The fix is not more argument. du/d(sin^2 gamma) = alpha^2 = 0.98 here, which is high, but sin^2(gamma) is already 0.883 and can rise at most to 1.000 — worth f(1.00) − f(0.8854) = 2.0794 − 1.9738 = 0.1056 in total. Whereas du/d(beta^2) = cos^2(gamma) = 0.117 looks small, but beta^2 can rise from 0.02 to anywhere. Rule 10.4 fires: buy the scarce source.

10.9 Why institutions systematically buy crossing and not divergence

Four structural reasons, each with a named counterpart elsewhere in this manual.

(1) Observability asymmetry. u_cross is a property of the roster; u_div is a property of behaviour over time, and only one is auditable from a spreadsheet. Star and Ruhleder apply exactly: functioning dissent is invisible, and its absence becomes visible only when the organization walks into a wall it had six people qualified to see.

(2) Goodhart capture. Once u_cross is a reported target, the cheapest way to raise the report is to relabel categories rather than acquire orthogonal ones (Chapter 8). u_div is harder to capture that way, which is precisely why nobody measures it: it produces uncomfortable numbers that administrative action alone cannot improve.

(3) The cost term punishes divergence directly. Internal dissent raises C_t immediately and visibly: meetings run longer, decisions take another cycle. Cross-hiring also raises C_t, but that increase is booked as onboarding and treated as investment. Same term, opposite political treatment.

(4) Conway's Law. An organization that resolves disputes by escalation builds systems with single points of arbitration; one running on shared vocabulary builds systems with one ontology. The architecture then reinforces the geometry that produced it, and gamma falls further.

A fifth is not structural but worth naming: **permitting divergence requires an unglamorous kind of managerial nerve, and it is the one lever of the two that cannot be delegated to a recruiting function.** Badged as observation, not result.

10.10 Instruments — measuring the two components separately

The decomposition is only useful if beta^2 and gamma can be estimated separately from real data. Both can. Every estimator below is badged: these are operational choices in the sense of Paper E §5.1, not consequences of the theorem.

Instrument A — the evaluability partition (estimates beta^2) For each ordered pair (x, y), sample y's recent work products and have x classify each item:

```
E  : "I can evaluate this on its merits."
N  : "I can read it but cannot judge whether it is right."
```

Then

```
beta^2_hat(x, y)  =  ( count of N ) / ( count of E + count of N )
```

Notes and cautions. The estimator is directional: beta^2(x,y) need not equal beta^2(y,x), and in hierarchical organizations it systematically does not. Use the mean for pair-level u_cross and *log the asymmetry separately* — it is itself a finding (§10.13). Classify against actual artifacts, not job titles: asking "can you evaluate an epidemiologist's work?" measures self-image, while handing someone three anonymized memos measures evaluability. And anonymize provenance, because a known author contaminates the E/N call with deference.

Instrument B — the shared-ground correlation (estimates gamma) Restrict to items both parties classified E. Have each score them independently and *before any discussion* — the ordering is not optional; see Chapters 11 and 22. Compute the correlation r of the two score vectors. Then

```
sin^2(gamma)_hat = 1 − r^2
u_div_hat        = alpha^2_hat ∗ (1 − r^2)  =  (1 − beta^2_hat)(1 − r^2)
```

Badge: 1 − r^2 is the standard coefficient of non-determination, and reading it as sin^2(gamma) is natural when the two score vectors stand for the state vectors restricted to the shared block. It is a modeling choice — defensible, the one this manual makes, and declarable in any preregistration (Chapter 25).

Cautions.

- r must be computed on *independent* scores. A correlation measured after discussion measures the discussion, not the participants. This single error accounts for most inflated agreement statistics in practice.
- Use enough items. Below about 12 shared-ground items the standard error on r swamps the estimate. Report the item count with every gamma.
- Negative r gives the same 1 − r^2 as +|r|: the estimator measures angle, not direction. **That is a genuine trap.** Two participants who systematically take opposite positions are *parallel* in this geometry (x . y = −1, u = 0), and the professional contrarian therefore buys nothing. If r sits near −1 rather than near 0, you do not have divergence, you have a mirror. See Failure Mode 10-F. Antiparallel is parallel.

Instrument C — the ensemble roll-up The theorem governs a pair; aggregating to a group is an operational choice.

```
u_cross_hat = mean over all ordered pairs of beta^2_hat
u_div_hat   = mean over all ordered pairs of alpha^2_hat (1 - r^2)
u_hat       = u_cross_hat + u_div_hat
```

Badge, and this matters: **the mean-over-pairs roll-up is a convenience, not a theorem** — a summary statistic, adequate for design and inadequate for a claim of proof. Chapter 24 requires the choice stated; Chapter 25 requires it frozen before fitting.

Instrument D — the disagreement log (a delta instrument) Levels are not enough: as Chapter 11 establishes, `u_div` decays. Keep a standing log:

```
For each substantive decision, record:
  - was a dissenting position stated on the record?  (Y/N)
  - by whom, and in which block?
  - from INSIDE the deciding block (u_div) or a different one (u_cross)?
  - was the dissent's substance addressed in the decision record?
```

The fourth field is the one that matters. Dissent recorded and not addressed is theatre and scores zero. Chapter 4's `eta` is the relevant term: **a disagreement that leaves no trace in the artifact converted none of its surplus into stored order, and `eta * S` for that interaction was zero regardless of how vigorous the meeting felt.**

10.11 Where the two components sit in the storage law

Nothing here changes the storage law. Both components feed the same scalar:

`Q_{t+1} = (1 - delta) Q_t + eta S_t - C_t,     S_t = f(u_cross,t + u_div,t)`

But the components have different *dynamics*, which is why they must be tracked separately even though the law only ever sees their sum.

`u_cross` decays slowly — set by the roster, moving at attrition timescales, so an annual recruiting and rotation review suffices. `u_div` decays fast, with a half-life in months (Chapter 11 computes it), and needs continuous protocol: independent-estimate-first, a dissent log, a red-team rotation.

The practical consequence for the headroom gauge (Chapter 6), `h_t = eta ln N - delta Q_t - C_t`: it does not contain u at all. u determines how much of the available headroom you convert, via `S_t = f(u) <= ln N`. **A collapse in `u_div` does not show up in the headroom gauge; it shows up as a persistent gap between `eta S_t` and `eta ln N` — between what the system could take in and what it does.** Define and track it:

```
UTILIZATION OF CEILING

  rho_ceil,t := S_t / ln N = f(u_t) / ln N,  in [0, 1]

  rho_ceil = 1   u = 1, orthogonal, running at the per-interaction ceiling
  rho_ceil = 0   u = 0, parallel, no surplus at any headcount
```

At `N = 6` (`ln 6 = 1.7918`): an organization at `u = 0.75` runs at `1.5581 / 1.7918 = 0.8696`; at `u = 1.00` it runs at 1.000. That 13 percent of unused ceiling is the entire prize in the worked example below, and it is invisible on every gauge except this one.

10.12 Semilattice, not tree

Alexander's "A City is Not a Tree" (1965) is the architectural statement of this chapter's structural claim. A tree partitions — every element in exactly one parent set, no overlap. A semilattice permits overlap.

A **tree** assigns each person exactly one block, so any pair is either in the same block (`beta^2 = 0`, all difference from `gamma`) or in different blocks (`beta^2 = 1`, `alpha^2 = 0`, so `u_div = 0` identically — *no shared ground* on which to disagree). **A strict tree forces every pair into a corner of the budget: pure crossing or pure divergence, never a portfolio**, every pair pinned and every balance index zero. A **semilattice** lets people hold partial membership in several blocks, so `0 < beta^2 < 1` for most pairs, both components are live, and B is positive across the ensemble.

This is why matrix structures, dual reporting, rotating secondments and shared instrumentation carry more surplus than the same headcount in a clean hierarchy — and why they cost more in `C`. Chapter 5's arithmetic and this chapter's are in tension by construction; Chapter 26 is where the trade is made explicitly.

10.13 The asymmetry finding

Instrument A's directional `beta^2` is diagnostic in its own right. Define `A(x,y) := beta^2_hat(x,y) - beta^2_hat(y,x)`. `A` near zero means mutual evaluability. Large positive `A` for a consistently-appearing `x` means *x reports being unable to evaluate others' work while others report being able to evaluate x's* — which has two very different causes. Either a genuine specialization gradient (x works deep and narrow), or **deference collapse**: `x` is junior and has learned not to claim evaluative standing, which is `u_div` being destroyed by rank and reappears in Chapter 17 as the head-of-table effect. Distinguish by anonymizing the artifacts. If `A` collapses when authorship is hidden, the asymmetry was deference. Cheapest decisive test in the chapter; it costs one afternoon.

The Specification

Requirements are numbered `SPEC 10.n` and each is checkable. A conforming system satisfies all of Group A, and satisfies Group B or records a dated exception with a named owner.

Group A — measurement (mandatory)

SPEC 10.1 The survey (Chapter 24) shall enumerate the `N` blocks in writing before any `u` estimate is computed, declared as *distinct methods or evaluative standards*, not job titles or demographic labels.

SPEC 10.2 `u_cross` and `u_div` shall be estimated and reported separately. A report containing only a total `u` is non-conforming.

SPEC 10.3 `beta^2` shall be estimated by the evaluability partition (Instrument A) against real, anonymized artifacts. Self-report against job titles does not conform.

SPEC 10.4 gamma shall be estimated from independent scores collected *before* any group discussion of the items. Post-discussion correlation is non-conforming and shall be labeled as such if reported.

SPEC 10.5 Every gamma estimate shall carry its shared-ground item count n_E. Estimates with n_E < 12 shall be marked provisional.

SPEC 10.6 The pair-to-ensemble roll-up rule shall be stated in the report and shall not change between periods without a dated note.

SPEC 10.7 The directional asymmetry A(x,y) shall be computed and reported. Any participant whose mean A exceeds 0.25 shall trigger the anonymization test of §10.13.

SPEC 10.8 The balance index B = [f(u) − max(f(u_cross), f(u_div))] / f(u) shall be computed each period and reported alongside u.

SPEC 10.9 The ceiling utilization rho_ceil = f(u) / ln N shall be computed each period and reported alongside h_t.

Group B — design (mandatory unless exception recorded)

SPEC 10.10 No plan shall use f(u_cross) + f(u_div) as an estimate of surplus. Plans shall use f(u_cross + u_div). The upper bound may appear only as a stated impossibility limit.

SPEC 10.11 No plan shall assert synergy beyond additivity. A benefit believed to exist beyond it shall be recorded as an uncredited expectation and excluded from the numbers.

SPEC 10.12 The first increment of u in any period shall come from whichever source is scarcer, per Rule 10.4, unless a written justification records why the more abundant source was chosen.

SPEC 10.13 Where B < 0.05 for two consecutive periods, the system shall be recorded as single-sourced and a remediation owner named. B < 0.05 at high u is not a healthy state; it is a concentrated one.

SPEC 10.14 Every substantive decision record shall carry a dissent field with four entries: stated (Y/N), by whom, source (cross or div), addressed in the artifact (Y/N). A blank field is non-conforming; "none" conforms and is itself data.

SPEC 10.15 Dissent recorded but not addressed in the artifact shall be scored eta = 0 for that interaction and shall not count toward u_div maintenance.

SPEC 10.16 Independent estimates shall precede discussion wherever u_div is a design input. The ordering is a hard requirement; reversing it destroys both the quantity being measured and the quantity being used.

SPEC 10.17 Onboarding, vocabulary standardization, and composite incentive metrics — the standard C-reduction moves — shall carry an explicit u_div impact estimate before adoption. Cutting C by spending u_div is a legitimate trade and an illegitimate accident.

SPEC 10.18 No participant shall be required to disagree. Divergence protocols shall create the *opportunity* and the *protection* to dissent, never an obligation: an intervention removing the ability to decline is out of scope for this manual (CANON §5.4). Compelled dissent is also useless — it produces low-information objections that inflate the log without moving r.

SPEC 10.19 Under a strict hierarchy, at least one crosscutting structure (rotation, community of practice, shared instrument, dual assignment) shall exist, so that some pairs have `0 < beta^2 < 1`. A pure tree pins every pair at a budget corner (§10.12).

SPEC 10.20 The contrarian check shall run each period: any pair with shared-ground `r` below `−0.6` shall be flagged as a possible mirror rather than a divergence, and re-examined. Antiparallel is parallel.

Group C — reporting

SPEC 10.21 The period report shall contain at minimum:

```
PERIOD: ______   N (declared blocks): ____
u_cross: ______ (n pairs ____)   u_div: ______ (mean n_E ____)   u: ______
N_eff = 1+(N−1)u: ______   S = f(u): ______   ln N: ______
rho_ceil = S/ln N: ______   B (balance index): ______
floor = max(f(uc),f(ud)): ______
ceiling = f(uc)+f(ud): ______   [impossibility limit only]
decisions with recorded dissent: ____/____
dissents addressed in artifact:  ____/____
mean directional asymmetry A: ______
```

SPEC 10.22 Any use of axiom S3 (the affine effective count) shall carry its badge in the same sentence. `N_eff` is a design choice, not a derived fact.

SPEC 10.23 The report shall state, verbatim: *the surplus functional as deployed is a lower bound on the ensemble's information, not the information itself.*

The Failure Modes

Each entry: name, clinical description, damaged parameter, signature, measurement, fix. These feed the catalog in Chapter 27.

10-A . THE FURNISHED MONOCULTURE

Clinical. High `u_cross`, `u_div` at or near zero. Every discipline present; each represented by exactly one permitted view.

Damages. `u_div` -> 0, pinning `f(u)` at `f(u_cross)` and driving `B` to zero. Robustness collapses while surplus still reads well.

Signature. Blank or "none" dissent fields across consecutive decisions; approval near 100 percent with substantive comments; "aligned" used as a success report; outsiders stop objecting within two quarters, and this is praised.

Measurement. Instrument B on within-block pairs: `r > 0.9`. `B < 0.05`. Fewer than one recorded-and-addressed dissent per ten decisions.

Fix. Rule 10.4: buy the scarce source. Independent-estimate-before-discussion (SPEC 10.16); a rotating, protected adversarial role; a non-blank dissent field (SPEC 10.14); an artifact recording how each dissent was addressed (SPEC 10.15). Expect `C` to rise; budget it against `h_t` first.

10-B . THE DEEP NARROW WELL

Clinical. High `u_div`, `u_cross` near zero. Vigorous argument inside a single block; no component of anyone's state outside it.

Damages. `u_cross` -> 0. Ashby violation: variety in one dimension, none in the others.

Signature. Reputation for rigour; failures always classified as "outside our remit"; recruitment selects for people who survive the internal debate, which selects for insiders.

Measurement. Instrument A: `beta^2` below 0.10 for nearly all pairs.

Fix. Import a practitioner from an unrelated trade with real standing, not an advisory seat; give them work products to evaluate, not presentations to attend. At large `gamma` the first cross-hire buys less than expected, since `du/d(beta^2) = cos^2(gamma)` is small. Buy anyway: the ceiling on further `gamma` is close, the ceiling on `beta^2` is not.

10-C . THE ADDED-BOUND PLAN

Clinical. A plan or forecast computing surplus as `f(u_cross) + f(u_div)`.

Damages. Nothing in the system; everything in the forecast. Overstates by 16 to 45 percent (§10.0.3), worst where the plan is most ambitious.

Signature. Two surplus numbers in a document, added. Language about the combination being "greater than the sum of its parts."

Measurement. Recompute with `f(u_cross + u_div)`. The gap is the overstatement.

Fix. SPEC 10.10 and 10.11. Retire the upper bound from planning use entirely.

10-D . THE C-FOR-DIVERGENCE TRADE MADE BY ACCIDENT

Clinical. An organization cuts coordination cost by standardizing vocabulary, induction, incentives and escalation, destroying `u_div` without deciding to.

Damages. `C` down (good), `u_div` down (unaccounted). The net effect on `eta S_t − delta Q_t − C_t` is often positive in the short run and negative once `S_t` responds, because the low-pass lag (Chapter 2) delays the `S` decline past the point where anyone attributes it to the change.

Signature. A cost-efficiency programme completed 6 to 18 months before an unexplained decline in decision quality — unsuspected, because it "worked."

Measurement. Track `u_div` and `C` on one time axis: a `C` step down followed by a `u_div` slope down. Chapter 25's shuffled-input control separates it from coincidence.

Fix. SPEC 10.17. Trade deliberately or not at all.

10-E . THE RELABELED BLOCK

Clinical. `N` raised on paper by splitting an existing category rather than adding an orthogonal one.

Damages. `N`, `beta^2`, and `u_cross` all inflated; every downstream number wrong in the optimistic direction.

Signature. Block boundaries following the reporting line rather than the method; blocks whose members trained in the same programme; `beta^2` near zero for pairs the declaration puts in different blocks.

Measurement. Cross-tabulate declared block against measured `beta^2`. If cross-block `beta^2` is not materially higher than within-block, the declaration is fiction.

Fix. Re-declare blocks by evaluative standard (SPEC 10.1). A block is defined by *what counts as a good answer*, not by who reports to whom.

10-F . THE MIRROR MISTAKEN FOR A WINDOW

Clinical. A participant who reliably opposes, credited as the source of `u_div`. In the squared geometry, systematic opposition is parallelism: `x . y = −1` gives `u = 0`.

Damages. `u_div` credited where none exists; worse, a designated opposer suppresses dissent from others, who assume the slot is filled.

Signature. One person objects to everything, and the objection carries no information because it is predictable from the proposal. `r` near −1.

Measurement. SPEC 10.20. Also: predict the objector's position from the proposal alone. If you can do it reliably, they are a mirror.

Fix. Rotate the adversarial role rather than appointing it (Chapter 9); score dissent by whether it changed the artifact. `r` near 0 is the target.

10-G . DEFERENCE COLLAPSE

Clinical. Directional asymmetry `A` driven by rank rather than specialization. Juniors report being unable to evaluate senior work; seniors report the reverse.

Damages. `beta^2` overestimated in one direction — silence recorded as crossing — and `u_div` destroyed, the party who would disagree having disclaimed standing to do so.

Signature. `A` correlates with seniority across the roster.

Measurement. §10.13's anonymization test.

Fix. Anonymize artifacts at the point of evaluation. Chapter 17's protocols — last-speaks-first for rank, silent-write-first, no head of table — attack the same fault from the geometry side.

10-H . THE UNADDRESSED DISSENT LOG

Clinical. Dissent recorded diligently and never changing an artifact. The log grows, `Q` does not.

Damages. `eta` -> 0 for the divergence channel: all the `C` of the argument, none of the storage.

Signature. A healthy dissent-stated rate and a near-zero dissent-addressed rate. Participants describe the process as going through the motions.

Measurement. The two counts in SPEC 10.21's report block, side by side. Their ratio is the divergence channel's `eta`.

Fix. SPEC 10.15. A Chapter 4 intervention, not a Chapter 10 one: the geometry was fine, the conversion was zero.

Worked Example

The Institute

A 34-person applied research institute working on flood risk, founded eleven years ago on an explicit cross-disciplinary premise and seriously executed.

Survey (Chapter 24 procedure). Declared blocks, by evaluative standard rather than by title:

```
N = 6                                       what counts as a good answer
  B1  hydrology / hydraulic modelling        model skill scores
  B2  civil & structural engineering         code compliance, FEA
  B3  economics & cost-benefit               discounted NPV, elasticities
  B4  public health & emergency response     outcome rates, response times
  B5  law & regulatory                       statutory defensibility
  B6  community engagement                   participation, equity metrics

ln N = ln 6 = 1.791759

Storage-law parameters (Chapters 3–5 procedures, prior eight quarters):
  delta = 0.08   quarterly loss of stored order without maintenance
  eta   = 0.45   fraction of surplus converted to durable structure
  C_t   = 0.20   quarterly dissipation: coordination, translation, compliance
  Q_0   = 6.00   current stored order, in the survey's Y-proxy units
```

Instrument A run. 1,122 ordered pairs, 5 anonymized artifacts per target, E/N classification. `u_cross_hat = mean beta^2 = 0.72` — real, and the institute earned it. Cross-block pairs show `beta^2` from 0.81 to 0.94 and within-block pairs 0.04 to 0.11, so the declared blocks survive the Failure Mode 10-E check.

Instrument B run. Mean `n_E = 19` shared-ground items per pair, above the provisional threshold; independent scores before discussion, 9-point scale, on 19 recent risk determinations.

```
within-block pairs:  mean r = 0.985  ->  1 - r^2 = 0.0298
cross-block pairs:   mean r = 0.94   ->  1 - r^2 = 0.1164 (but alpha^2 is small)

ensemble roll-up:  u_div_hat = mean over pairs of alpha^2 (1 - r^2) = 0.03
```

The as-found decomposition.

```
u_cross = 0.72
u_div   = 0.03
u       = 0.75

f(u_cross) = ln(1 + 5*0.72) = ln(4.60)  = 1.526056
f(u_div)   = ln(1 + 5*0.03) = ln(1.15)  = 0.139762
f(u)       = ln(1 + 5*0.75) = ln(4.75)  = 1.558145

floor   = max(1.526056, 0.139762) = 1.526056
ceiling = 1.526056 + 0.139762     = 1.665818   [impossibility limit only]

N_eff   = 1 + 5*0.75 = 4.750       (S3 is a design axiom, badged)
rho_ceil = 1.558145 / 1.791759 = 0.8696
B       = (1.558145 - 1.526056) / 1.558145 = 0.0206
```

Read the diagnosis. `B = 0.0206`, below the SPEC 10.13 threshold: the institute is single-sourced, `u_div` contributing 2.06 percent of its surplus. **Thirty-four people, six genuine disciplines, and a surplus that would be almost unchanged if every within-block disagreement in the building were silenced tomorrow.** The dissent log over the last 40 decisions:

```
decisions with dissent stated:              6 / 40
  of which from a different block (cross):  6
  of which from within the deciding block:  0
dissents addressed in the decision artifact: 2 / 6
mean directional asymmetry A:              0.31   (exceeds SPEC 10.7 trigger)
```

Zero within-block dissents in forty decisions. Not low. Zero. And `A = 0.31` triggers the anonymization test, which the institute ran: `A` fell to 0.09 with authorship hidden. **Deference collapse (Failure Mode 10-G) is present and is part of why `u_div` is near zero.**

Headroom check before proposing anything (Chapter 6), and current behaviour.

```
h_0 = eta ln N - delta Q_0 - C_0
    = 0.45*1.791759 - 0.08*6.00 - 0.20
    = 0.806292 - 0.480000 - 0.200000 = 0.126292

S_0     = f(0.75) = 1.558145
eta S_0 = 0.45 * 1.558145 = 0.701165
Delta Q = 0.701165 - 0.680000 = +0.021165   -> growing, barely
Q_1     = 0.92*6.00 + 0.701165 - 0.20 = 6.021165
Q_inf   = (0.701165 - 0.20) / 0.08 = 6.264563
```

`h_0 > 0`: growth is available, and `r_max = h_0 = 0.126292` (Paper E, Cor 4.2) — no policy based on interaction geometry alone can guarantee more. The institute grows at 0.021 per quarter toward 6.265, visibly flat:

```
t:    0      1      2      3      4      5      6      8     10     12
Q:  6.000  6.021  6.041  6.059  6.075  6.090  6.104  6.129  6.150  6.167   -
> 6.265
```

The wrong intervention. The standing proposal is a seventh block — remote sensing — at a cost of two senior hires. Before pricing it, run the `u_cross` lever within the existing `N`, holding `u_div`

= 0.03:

```
u_cross   u      S = f(u)    Delta Q = etaS − deltaQ_0 − C
-------------------------------------------------------------
 0.60    0.63    1.423108      −0.039601
 0.72    0.75    1.558145      +0.021165     <− as found
 0.85    0.88    1.686399      +0.078880
 0.97    1.00    1.791759      +0.126292     <− at the ceiling
```

Driving u_cross from 0.72 to 0.97 — an enormous restructuring — buys +0.126292, exactly h_0: the whole headroom, at the highest possible price.

The right intervention. Hold u_cross fixed. Buy the scarce source.

The budget constraint (§10.3) caps u_div at alpha^2 = 0.28 while u_cross = 0.72, and reaching that cap requires within-block r = 0 on shared ground — full independence between two people in the same discipline. The institute will not hit it. The arithmetic of partial achievement is what matters:

holding u_cross = 0.72, N = 6, delta Q_0 + C = 0.68:

```
u_div    u      S = f(u)    gain in S   Delta Q per quarter
-------------------------------------------------------------
 0.00   0.72    1.526056      +0.0000      +0.006725
 0.03   0.75    1.558145      +0.0321      +0.021165     <− as found
 0.05   0.77    1.578979      +0.0529      +0.030540
 0.10   0.82    1.629241      +0.1032      +0.053158
 0.15   0.87    1.677097      +0.1510      +0.074693
 0.20   0.92    1.722767      +0.1967      +0.095245
 0.28   1.00    1.791759      +0.2657      +0.126292
```

Target: protocols moving within-block r from 0.985 to roughly 0.71, giving 1 − r^2 = 0.50 and an ensemble u_div at the full remaining budget of 0.28. Ambitious; a realistic first year is half of it.

Full target (u_div = 0.28, u = 1.00), against the as-found state:

```
                    S        eta S     Delta Q at t=0     Q_inf
as-found         1.558145   0.701165    +0.021165        6.264563
u_div raised     1.791759   0.806292    +0.126292        7.578650
                                        6.0x faster      +1.314087  (+21.0%)

rho_ceil goes 0.8696 -> 1.000.
Delta Q = 0.126292 = h_0 exactly: the intervention consumes the entire
available headroom and not one unit more. Corollary 4.2 doing its job --
r_max = h_t, and running the geometry to u = 1 is the only way to reach it.
```

Twenty-one percent more stored order at steady state, and a six-fold increase in the quarterly growth rate, purchased entirely by permitting people already on the payroll to disagree with colleagues in their own discipline. No hires, no new block, no change to u_cross, delta or eta — except that C will rise, which is the next paragraph.

Cost honesty. Structured within-block dissent raises C — longer decision cycles, one extra review pass, facilitation time. Estimate +0.03 per quarter:

```
C_1 = 0.23
h_1 = 0.806292 − 0.08∗6.00 − 0.23 = 0.096292      (was 0.126292)
Delta Q at u=1.00 with C=0.23 = +0.096292
Q_inf = (0.806292 − 0.23) / 0.08 = 7.203650
```

Still 7.204 against 6.265: a gain of 0.939, or 15.0 percent. **The intervention survives its own cost by a wide margin, because the surplus gain is bought at the ceiling of the u parameter, where the institute had 13 percent of its per-interaction capacity sitting unused.**

Now the seventh-block proposal under the same cost honesty. It raises `eta ln N` from 0.806292 to `0.45 ∗ ln 7 = 0.875660`, but two senior hires raise C by an estimated 0.09 (salary is not in C; translation burden, extra coordination paths per Brooks, and onboarding are). At the as-found `u = 0.75`, now with `N = 7`:

```
S = ln(1 + 6∗0.75) = ln(5.50) = 1.704748
eta S = 0.767137
Delta Q = 0.767137 − 0.480 − 0.29 = −0.002863     -> declining
Q_inf = (0.767137 − 0.29) / 0.08 = 5.964213      -> BELOW the as-found 6.265
```

Adding a seventh discipline to a system with `u_div = 0.03` makes it worse. The ceiling rose, the cost rose, the geometry did not. This is Chapter 8's warning in its sharpest form: you cannot escape a headroom problem by adding domains.

The specification the institute adopts.

```
INTERVENTION SPEC — INSTITUTE, QUARTERS 1–4

1. Independent-estimate-before-discussion on every risk determination; scores
   to a locked form that closes before the meeting opens. (SPEC 10.16)
   Owner: research director.

2. Dissent field mandatory on all decision records, blank not permitted;
    artifact states each dissent's resolution as accepted / rejected-with-
reason
   / deferred-with-date. (SPEC 10.14, 10.15) Owner: secretariat, chair.

3. Artifacts anonymized at peer evaluation, per the deference finding
   (A fell 0.31 -> 0.09 under anonymization).

4. Rotating within-block adversarial assignment, one quarter per person,
   protected and voluntary; declining carries no penalty and is recorded as a
   decline, not a failure. (SPEC 10.18) Owner: block leads.

5. Quarterly re-run of Instruments A and B, reported per SPEC 10.21. C
   monitored monthly; if C exceeds 0.26, pause item 4 and re-cost, because
   h_t must not go below zero (Chapter 6, Cor 3.4).
```

Falsification criteria (Chapter 25 form), preregistered and frozen.

```
PREDICTION: within-block r on shared ground falls from 0.985 to below 0.85
            within four quarters, giving u_div >= 0.10 and S >= 1.629.

FALSIFIED IF, at quarter 4:
  (a) r remains above 0.95, OR
  (b) u_div < 0.06 despite full protocol compliance, OR
  (c) Y (the order proxy) has not increased relative to the frozen
      persistence baseline, OR
  (d) the shuffled-input control fits as well as the true input series.

REQUIRED BASELINES: persistence; direct-input; matched autoregressive;
                    shuffled-input control.

NOTE ON (c): the order proxy Y here is the count of adopted, maintained
technical standards in force — chosen because it is NOT a repackaging of U.
It contains no measure of disagreement, interaction volume, or staffing
composition. (Paper E, §6.3.)
```

What the numbers do not say. `u_cross = 0.72` is real, was expensive, and remains the larger half of the institute's surplus by a factor of eleven. Nothing here argues against cross-disciplinary staffing. The finding is narrower and harder: **the institute completed one of two purchases, declared the programme finished, and left the cheaper half of its own portfolio on the table for eleven years — because the half it bought is visible on a spreadsheet and the half it did not is visible only in behaviour nobody was recording.**

That is a measurement failure before it is a design failure. Chapter 11 shows what happens when an organization does not merely fail to buy `u_div` but actively loses the `u_div` it started with — and why that makes difference a maintenance obligation rather than an acquisition.

Chapter 11 . Pseudo-Diversity and the Parallel Wall

This is the failure mode that costs modern institutions the most, and it is almost never diagnosed, because every instrument the institution owns reports that nothing is wrong.

The parallel wall is not a shortage of people, talent, funding, effort, or good intention. It is a geometry in which many capable people, each doing competent work, collectively return the surplus of one. The arithmetic that prices it is short enough to fit on a card, and this chapter's first job is to put it there.

Its second job is harder. u is not a property an organization acquires and then owns. **It is a quantity that decays, and that makes non-parallelism a maintenance obligation, not a procurement one.** A group can be assembled correctly, measure well on day one, and be parallel within two years without a single person leaving, a single policy changing, or anyone making a decision they would recognize as wrong.

The Law

11.0.1 The arithmetic of zero

From Paper A, Theorem 2.1 and the effective count that follows from axiom S3:

```
N_eff = g(u) = 1 + (N−1)u
S     = F(u) = f(u) = ln(1 + (N−1)u) = ln N_eff

at u = 0:   N_eff = 1 + (N−1)(0) = 1
            S     = ln 1 = 0
```

That is the whole indictment. **Three glasses ground from one doctrine are one glass, and their agreement is not evidence of anything except that they were ground from one doctrine.**

Paper A Theorem 4.3 (extremality) makes it exact and two-sided: `F = 0` if and only if the states are parallel, and `F = ln N` if and only if they are orthogonal. There is no gradation at the bottom. Parallelism does not produce a small surplus; it produces none.

Note also what the law does *not* say. It does not say the group produces nothing. Each member still does work. What returns zero is the *interaction surplus* — the part of the output that exists because there is more than one of them. **A parallel group of forty is not a bad group. It is a group of one, paid forty times, and audited as though it were forty.**

11.0.2 What "parallel" means operationally

`u = 1 − (x · y)^2`. Two states are parallel when their normalized inner product has magnitude 1 — when one is a scalar multiple of the other. In an organization the operational readings are:

Formal	Operational reading
`x · y = +1`	the second party's judgment is predictable from the first's, with the same sign

Formal	Operational reading
`x . y = −1`	the second party's judgment is predictable from the first's, with the sign flipped
`u = 0`	in either case: knowing the first tells you the second, so the second carried no new information

The second row is the trap Chapter 10 named (Failure Mode 10-F) and it bears repeating here because pseudo-diversity programmes frequently buy it by mistake. **Antiparallel is parallel.** A designated opposer whose position is derivable from the proposal is a mirror, not a window, and the law prices a mirror at zero.

11.0.3 The price of pseudo-diversity, in numbers

Take a declared roster of N categories. Compute what the roster is actually worth at various levels of measured non-parallelism.

```
      N = 5              N = 10             N = 20             N = 40
 u    N_eff     S       N_eff     S        N_eff     S        N_eff     S
-------------------------------------------------------------------------
---
0.00   1.000   0.0000    1.000   0.0000     1.000   0.0000     1.000   0.0000
0.01   1.040   0.0392    1.090   0.0862     1.190   0.1740     1.390   0.3293
0.02   1.080   0.0770    1.180   0.1655     1.380   0.3221     1.780   0.5766
0.05   1.200   0.1823    1.450   0.3716     1.950   0.6678     2.950   1.0818
0.10   1.400   0.3365    1.900   0.6419     2.900   1.0647     4.900   1.5892
0.20   1.800   0.5878    2.800   1.0296     4.800   1.5686     8.800   2.1748
0.40   2.600   0.9555    4.600   1.5261     8.600   2.1518    16.600   2.8094
1.00   5.000   1.6094   10.000   2.3026    20.000   2.9957    40.000   3.6889
```

Read the `N = 40` column. A forty-category roster at `u = 0.05` has an effective count of 2.95. **Forty declared categories, three effective ones — and the organization is paying for forty, sizing its meeting rooms for forty, and carrying the coordination cost of forty.** Chapter 5's arithmetic completes the picture: at `N = 40`, Brooks' communication paths number `n(n−1)/2 = 780`. The group carries 780 channels of cost to deliver the surplus of three.

That ratio — cost of the declared count, surplus of the effective count — is the entire economics of the parallel wall, and it is why pseudo-diversity is expensive rather than merely disappointing.

11.0.4 The steepest slope is at zero

Paper A Theorem 4.2 gives the sharp Lipschitz constant `N − 1`. Differentiating:

```
f'(u) = (N−1) / (1 + (N−1)u)

f'(0)  = N − 1        <-- the maximum, attained exactly at parallelism
f'(1)  = (N−1)/N      <-- the minimum
```

The first unit of difference is the most valuable unit of difference there will ever be. Concretely:

N	f(0.01) − f(0)	f(1.00) − f(0.99)	ratio
3	0.0198	0.006689	3.0
5	0.0392	0.008032	4.9
10	0.0862	0.009041	9.5
20	0.1740	0.009545	18.2
40	0.3293	0.009798	33.6

At N = 40, the first hundredth of u is worth 33.6 times the last hundredth. At N = 10, moving u from 0.00 to 0.05 buys 0.3716; moving it from 0.95 to 1.00 buys 0.0460 — a factor of 8.07.

This cuts both ways, and the second cut is the one this chapter is about. **Because the function is steepest at zero, a group near the parallel wall loses more surplus per unit of u lost than a group anywhere else on the curve.** The last increments of difference to disappear are the most expensive ones. Decay does not slow as it approaches zero; in surplus terms it accelerates.

11.0.5 u as a decaying state variable

Nothing in Paper A makes u a function of time. Paper A is static geometry. The storage law (Paper E §3) evolves Q, and takes U_t as an *input*:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1)U_t)
```

U_t is exogenous to the law. **Nothing in the mathematics says U_t stays where you put it, and everything in the empirical literature says it does not.**

This manual therefore adopts an explicit modeling postulate, badged here and carried through the chapter:

> **POSTULATE 11.1 (u-decay), badged.** In a group under ordinary communication, absent maintenance, measured non-parallelism decays geometrically: U_{t+1} = (1 − lambda) U_t with lambda in [0,1).
>
> This is a modeling assumption, not a theorem. It is not in Paper A, not in Paper E, and not proved anywhere in this book. It is adopted because it is the simplest form consistent with the cascade mechanism of §11.4, because it makes lambda estimable from two measurements, and because it is falsifiable in the manner of Chapter 25. If measured u under zero maintenance does not decline approximately geometrically, replace the postulate.

With that badge in place, the structural claim of this chapter becomes sayable in the manual's own terms:

```
delta  is the decay rate of STORED ORDER      (Chapter 3)
lambda is the decay rate of DIFFERENCE        (this chapter)

Both are leaks. Both require maintenance. Only one of them is
routinely measured, and it is not lambda.
```

u is a delta problem as much as a design problem, and an organization that has a maintenance regime for its documents and none for its disagreement has budgeted for one leak and ignored the other.

11.0.6 Badges and honesty clauses

- **Proven (Paper A, Thm 2.1, 4.3):** `f(u) = ln(1 + (N−1)u)` under S1-S4; `F = 0` iff parallel, `F = ln N` iff orthogonal.
- **Proven (Paper A, Thm 4.2):** sharp Lipschitz constant `N − 1`, attained at `u = 0`.
- **Design axiom, not derived (S3):** the affine effective count `N_eff = 1 + (N−1)u`. Every `N_eff` in this chapter inherits that choice.
- **Postulated and badged (11.1):** geometric decay of `u`.
- **Honesty clause:** as deployed, the surplus functional is a **lower bound** on the ensemble's information, not the information itself. A group measured at `u = 0` on the observables you chose may hold information those observables do not see. It does not hold information the instrument *credits*, and stored order responds to what the instrument credits. **If your instrument reads zero and you believe the group is not parallel, the correct response is to change the instrument, not to overrule it.**
- **Honesty clause:** Paper A is pure mathematics and makes no empirical or metaphysical claim. Paper D: surplus is not identified with Shannon entropy or mutual information in general.
- **Refused (CANON §5.4):** nothing in this chapter recommends compelling anyone to hold, express, or perform a view. Every instrument here measures; every intervention here creates opportunity and protection. A person who declines to dissent has exercised a right, and their decline is recorded as data, not as a failure.
- **Scope note, stated once and meant.** This chapter treats diversity as an engineering parameter of an information-processing system. It makes no claim about the social, moral, or legal case for any staffing policy, and it takes no position on any of them. A reader who imports a political argument into this chapter has imported it; it is not here. **The fault named here is that a measurable quantity is going unmeasured, and the fix is to measure it.**

The Engineering

11.1 Three glasses ground from one doctrine

Three optical instruments are pointed at the same object. They agree, and the agreement is reported as convergent evidence.

But all three were ground on the same machine, to the same prescription, by the same house, and calibrated against the same reference. Their agreement is a statement about the machine. It contains exactly as much information about the object as one instrument does, and the confidence gained from the "three" is manufactured.

In the law: `x1 = x2 = x3` up to scale, so `u = 0` for every pair, `N_eff = 1`, `S = ln 1 = 0`. **The confidence went up and the information did not, and that gap is what pseudo-diversity produces as its characteristic output.**

Condorcet's jury theorem says the same from the other side: collective accuracy exceeds individual accuracy *only under independence*. Independence is not a nicety attached to the theorem, it is the theorem's engine, and dropping it drops the conclusion. §11.9 computes how fast.

11.2 The five species

Pseudo-diversity is not one fault but five, needing different instruments to detect and different interventions to fix.

Species 1 — declared variety with epistemic uniformity. The roster varies along one or more declared axes; the *evaluative standards* do not. Every member was trained to recognize the same things as good answers. The blocks were declared along an axis that is not the axis on which judgment varies, so `u_cross` measures high on the declared axis and near zero on the axis that governs the decisions being made.

Detection: declare blocks by evaluative standard (Chapter 10, SPEC 10.1), then re-measure. If `beta^2` does not rise across the declared boundary, the boundary is not an epistemic boundary.

Species 2 — the echo chamber. Information intake is filtered through a common upstream source. Members differ in training and converge in evidence, because they read the same five outlets, attend the same three conferences, and follow the same feed. Difference in the members does not survive identical inputs.

Detection: audit the information supply chain, not the roster. Ask each member independently for the three sources that most changed their view in the past year; compute the overlap. High overlap is the signature.

Species 3 — correlated forecasts. Members produce independent-looking numerical judgments whose errors correlate. Each estimate is honestly made; the ensemble is nonetheless one estimate with noise. This species does most damage in risk, finance, planning, and engineering safety, because the aggregate is *reported* as an ensemble and *behaves* as a point.

Detection: error-decorrelation measurement, §11.10. This is the one species with a clean, cheap, quantitative instrument, and it is still rarely run.

Species 4 — committee theatre. Structure is elaborate — subcommittees, review stages, sign-offs, quorums — and the geometry flat, because every stage draws from the same pool and applies the same standard. Each stage adds `C` and no `u`: the process grows quadratically in cost and not at all in surplus.

Detection: count the stages, then count how often a later stage reversed an earlier one on substance. A reversal rate near zero across many decisions means the later stages are re-reading the first stage's instrument.

Species 5 — the common training pipeline. Reviewer, candidate, and appointment pools all drawn from a few upstream institutions. This is Barabási-Albert preferential attachment operating on careers: hubs attract, hubs train, and everyone downstream carries the hub's prior. The structure is efficient, robust to random loss, and — per Albert, Jeong and Barabási — fragile to anything corrupting the hub, since everything downstream inherits it at once.

Detection: trace training lineage rather than affiliation. Two people at different institutions who trained under the same advisor are one lineage.

11.2.1 Manufactured opposition — a sixth fault, and not the Parallel Wall

The five species above all describe bodies that agree too much. This one describes a body that disagrees on schedule, and it was catalogued after the five because it does not look like them

from outside. It is entered here as a distinct fault, with its own signature and its own fix, and it is recorded as failure mode 11-I below.

The arithmetic is §11.0.2's second row, taken seriously. `u = 1 − (x . y)^2` depends only on `|rho|`, so **`rho = −1` returns `u = 0` exactly as `rho = +1` does.** Perfect opposition is parallel. The contrarian who reliably inverts you carries the same information as the sycophant who reliably echoes you: none, because in both cases their position is a function of yours.

"Creative tension" at `rho` near -1 is therefore not a weak purchase of `u`. It is a zero purchase of `u` with a cost overhead, and that overhead is what makes it distinct from the Parallel Wall rather than a variant of it. A parallel roster wastes seats it already had. A designated-opposer seat is *added* — hired, seconded, or scheduled — to supply a divergence it cannot supply. Price the seat alone, at `eta = 0.35`, `delta = 0.10`, and a seat cost of 0.05 per period:

the added seat	u	S	eta∗S	its Q_inf contribution
genuinely independent, `rho = 0`	1.0000	0.6931	0.2426	**+1.9260**
partially correlated, `rho = 0.6`	0.6400	0.4947	0.1731	+1.2314
perfect echo, `rho = +1`	0.0000	0.0000	0.0000	**-0.5000**
perfect opposition, `rho = −1`	0.0000	0.0000	0.0000	**-0.5000**

The last two rows are the same row. The opposer and the echo have identical consequences for stored order, and both are strictly worse than not filling the seat, because `C` is paid and nothing is returned.

The fault is worse than its arithmetic, for a reason that is not in the equation and is badged as observation rather than theorem: **the slot appears filled.** A body with a house sceptic believes it has bought dissent and stops looking for the independent judgment it actually needs, so the manufactured opposition suppresses demand for the real thing. That is the same structure as Chapter 9's captured red team (FM-9.3), arrived at from the opposite direction.

Detection: predict the objector's position from the proposal before the meeting. If the prediction is reliable, `|rho|` is near 1 and the seat is a mirror regardless of how heated the exchange is. Measure `|rho|`, never `rho`.

Fix: target `|rho|` near zero, not `rho` near -1. Rotate the role so it cannot become a stable inversion; score the seat by artifact change rather than by objection count. Chapter 7 §7.6 gives the geometry — the angle table folded about 90 degrees — Chapter 9 §9.2 gives the dyadic table, and Chapter 46 treats the two-person case in full.

11.3 The inference error, stated formally

Underneath all five species is a single reasoning mistake, worth stating in a form precise enough to check.

Agreement between two observers is evidence about the world **only to the extent that the observers could have disagreed.** The informativeness of concordance is a function of the prior probability of discordance; when that prior is zero, concordance carries no information.

```
CONCORDANCE TEST

  Before reading the second opinion, write down:
    p = P(this reviewer disagrees with the first, given the first)

  If p is near 0, the second opinion is not evidence, whatever it says.
  If you cannot estimate p, you cannot interpret the agreement.
```

This test costs thirty seconds and can be applied to any second opinion, any peer review, any confirming analysis. Most institutions have never asked it about any of theirs. **The question is not "do they agree?" but "could they have disagreed?" — and the second question is the only one whose answer carries information.**

Chapter 9's line is the same statement in the roster's language: a second opinion from someone whose training and incentives replicate the first is worth `ln(1) = 0`.

11.4 The cascade: how a diverse group becomes parallel

Everything so far describes a group that was parallel to begin with. The harder and more common case is a group that was genuinely diverse and became parallel while doing nothing wrong.

The mechanism is the information cascade of **Bikhchandani, Hirshleifer and Welch (1992)**, worth working through because it shows the collapse occurring through *individually rational* behaviour. No one defects, no one is coerced, and the group still ends up as one instrument.

The canonical setup: agents act in sequence on a binary question. Each holds a private signal, correct with probability `p > 1/2`. Each observes the actions of those who acted before, but not their private signals. Each acts to maximize their own accuracy.

Agent 1 has only their signal, so acts on it. Agent 2, seeing one action and holding one signal, acts on their own signal (breaking ties toward it). Agent 3 now faces two prior actions. If both agree, two units of evidence stand against Agent 3's one, and the rational act is to follow the crowd **and ignore their own signal.** From that moment, every subsequent agent sees only actions that carried no private information, so every subsequent action carries none either. The cascade is locked.

The arithmetic of the canonical model:

```
State A is true; signal accuracy p.
Agents 1 and 2 both follow their own signals.
  both correct   -> p^2          -> correct cascade
  both incorrect -> (1-p)^2      -> INCORRECT cascade
  mixed          -> 2p(1-p)      -> no cascade; the process restarts
```

p	P(correct cascade)	P(incorrect cascade)	E[agents whose signal entered the record]
0.55	0.5990	0.4010	3.96
0.60	0.6923	0.3077	3.85
0.70	0.8448	0.1552	3.45
0.80	0.9412	0.0588	2.94
0.90	0.9878	0.0122	2.44

Two results deserve to be read slowly.

First, the incorrect-cascade probability. With individually 70-percent-accurate agents, the group locks onto the wrong answer 15.5 percent of the time, and once locked, adding more agents does not help — every additional agent copies. **A room of forty 70-percent-accurate people has a roughly one-in-six chance of being confidently, permanently wrong, and the confidence grows with the room size while the accuracy does not.**

Second, the informative-agent count. At `p = 0.70`, on average 3.45 agents' private signals ever enter the public record. Everyone after that is echo. In a room of forty, thirty-six and a half people are re-reporting three and a half. The declared `N` is 40. The information-bearing count is 3.45.

That is pseudo-diversity manufactured in real time, from a group that was genuinely diverse when it walked in.

11.5 Cascades collapse u, and that is the delta problem

The cascade result is normally taught as a decision-theory curiosity. In this manual it is the mechanism behind Postulate 11.1, and the bridge is direct.

Before the cascade, participants' reported states differ because each reflects an independent private signal: `u > 0`. After the cascade, every reported state is a copy of the public history: `x_i = x_j` for all post-cascade `i, j`, so `u = 0` among them. Not small. Zero, by Theorem 4.3.

The key structural point is that **cascades do not need a new event to occur. They occur as a by-product of ordinary sequential communication, which is to say they occur continuously in any group that talks.** Every meeting in which the senior person speaks first, every review in which the second reviewer sees the first reviewer's score, every forecast round in which the group discusses before estimating, is a cascade opportunity. Most organizations present dozens per week.

That is why u decays rather than merely being low. And it is why the right frame is maintenance:

	acquisition problem	maintenance problem
question	"did we hire difference?"	"are we still different?"
measured	at recruitment	every period, forever
owner	recruiting	whoever runs the meetings
failure looks like	a homogeneous roster	a heterogeneous roster that produces homogeneous output

Almost every organization has the left column staffed and the right column unstaffed. **The left column is the one that photographs well.**

11.6 The decay table

Under Postulate 11.1, `u_t = u_0 (1 − lambda)^t`, and the half-life is

`t_half = ln 2 / ( −ln(1 − lambda) )`

which is the same formula Chapter 3 gives for `delta`, because it is the same kind of leak.

lambda	half-life (periods)	u after 10 periods, from u_0 = 0.60
0.02	34.31	0.4902
0.05	13.51	0.3592
0.10	6.58	0.2092
0.15	4.27	0.1181
0.20	3.11	0.0644
0.30	1.94	0.0169
0.50	1.00	0.0006

And the surplus consequence, at `N = 10`, `u_0 = 0.60`:

`lambda = 0.15,  N = 10`

t	u	N_eff	S = f(u)	S as % of ln 10
0	0.6000	6.400	1.8563	80.6%
2	0.4335	4.902	1.5895	69.0%
4	0.3132	3.819	1.3399	58.2%
6	0.2263	3.037	1.1107	48.2%
8	0.1635	2.471	0.9048	39.3%
10	0.1181	2.063	0.7242	31.5%
12	0.0853	1.768	0.5699	24.8%

Twelve periods — three years at quarterly cadence — takes a group from 80.6 percent of its per-interaction ceiling to 24.8 percent, with the same people, the same budget, and the same org chart. **Nothing in this table requires anyone to do anything wrong. It requires only that nobody does anything to stop it.**

Estimating lambda in a real system. Measure u at two times separated by k periods, with no u-maintenance intervention between them:

`lambda_hat = 1 − (u_1 / u_0)^(1/k)`

Two cautions, both mandatory. First, the interval must genuinely contain no maintenance, or you have measured the net of decay and repair and cannot separate them. Second, u must be measured the same way both times — the same instrument, the same item pool, the same pre-discussion ordering (Chapter 10, SPEC 10.4). A change of instrument between the two measurements produces a `lambda` that is an artifact of the change.

11.7 Why the institution looks healthy after its input has collapsed

Here is the property that makes this failure mode so hard to catch in time, and it comes directly from Chapter 2's circuit identity.

The storage law is a first-order low-pass filter. `Q` responds to `S` with a lag of order `1/delta` periods. So when `u` — and therefore `S` — starts falling, `Q` does not immediately fall. It keeps rising for a while, on the momentum of the accumulated term, and then rolls over.

Worked, with `N = 5`, `eta = 0.55`, `delta = 0.09`, `C = 0.15`, `Q_0 = 4.00`, `u_0 = 0.55`, and `lambda = 0.18` (half-life 3.49 periods):

t	u	N_eff	S	eta S	delta Q	Delta Q	Q	
0	0.5500	3.200	1.1632	0.6397	0.3600	+0.1297	4.0000	
1	0.4510	2.804	1.0310	0.5671	0.3717	+0.0454	4.1297	
2	0.3698	2.479	0.9080	0.4994	0.3758	−0.0264	4.1751	<– peak
3	0.3033	2.213	0.7944	0.4369	0.3734	−0.0865	4.1488	
4	0.2487	1.995	0.6905	0.3798	0.3656	−0.1358	4.0623	
6	0.1672	1.669	0.5121	0.2817	0.3376	−0.2059	3.7511	
8	0.1124	1.450	0.3714	0.2043	0.2985	−0.2442	3.3166	
10	0.0756	1.302	0.2642	0.1453	0.2537	−0.2584	2.8185	
12	0.0508	1.203	0.1851	0.1018	0.2071	−0.2553	2.3015	
16	0.0230	1.092	0.0879	0.0484	0.1192	−0.2208	1.3240	
20	0.0104	1.042	0.0407	0.0224	0.0460	−0.1736	0.5110	
23	0.0057	1.023	0.0227	0.0125	0.0022	−0.1398	0.0248	

`u` peaks at `t = 0`. `Q` peaks at `t = 2`. **For two full periods after the input has begun to collapse, the institution's stored order is still rising, and every dashboard says the programme is working.** By the time `Q` visibly falls, `u` has already halved.

That lag is not a defect of the model. It is the correct behaviour of an accumulator, and it is the reason the input must be instrumented directly. **Watching `Q` to detect a `u` collapse is watching the reservoir to detect the drought.**

Two further readings from this run.

The headroom gauge does not see it. Compute `h_t = eta ln N − delta Q_t − C_t` along the same path: `h_0 = 0.3752`, `h_2 = 0.3594`, `h_8 = 0.4367`, `h_16 = 0.6160`. Headroom *rises* as the institution collapses, because `Q` is falling and `delta Q_t` shrinks with it. Chapter 6's gauge is doing its job — it reports available capacity, and a dying institution has plenty. **`h_t` measures what you could take in; `rho_ceil = S_t / ln N` measures what you do. Only the second one detects a parallel wall.** Here `rho_ceil` goes 0.723 at `t = 0` to 0.055 at `t = 16`.

The model names its own breakdown. Admissibility requires `C_t <= (1 − delta)Q_t + eta S_t`. At `t = 23`, the right side is 0.0350 against `C = 0.15`. Admissibility fails, and beyond that the linear model reports negative `Q`, which is not a physical state. Read it correctly: the arithmetic does not predict a negative institution. **It names the period at which the institution can no longer pay its own running cost, and stops.** Paper E, Theorem 3.5 makes the same statement in equilibrium form: with `eta F(u*) < c*` the steady state is negative, and the theorem does

not hide the failure mode. At `u = 0.08` here, `Q_inf = (0.55 * 0.277632 - 0.15)/0.09 = 0.030` — the regime is at the edge of unsustainable, and a little below that u it is over it.

11.8 Detection instruments

Five instruments, in increasing order of cost. Run at least two; they fail differently.

Instrument 1 — the collision test (Book One's method, formalized). Put a live question to two members separately, in writing, with no prior discussion. Compare. Where the two accounts *collide* — genuinely conflict, not merely differ in emphasis — you have non-parallelism. Where they read as one account in two hands, you do not.

Score it: over `m` questions, the collision rate `c = collisions / m` is a crude but honest first estimate of the group's u on that question class. It is crude because it binarizes a continuous angle. It is honest because it cannot be faked by anyone who has not seen both answers.

Instrument 2 — the disagreement audit. Retrospective, cheap, and usually devastating. Pull the last 40 substantive decision records. Count:

```
decisions with any dissent recorded                    ____ / 40
dissents from inside the deciding discipline (u_div)   ____
dissents from a different discipline (u_cross)         ____
dissents whose substance changed the final artifact    ____
reversals of an earlier stage by a later stage         ____
```

The fourth line is the one that matters, per Chapter 4: dissent that leaves no trace in the artifact had `eta = 0` and converted nothing. The fifth line detects Species 4.

Instrument 3 — error-decorrelation measurement. The strongest instrument available, and the subject of §11.10. Requires a history of independent, resolvable predictions. Where it can be run, run it; it produces a number rather than an impression.

Instrument 4 — the pre-mortem. Prospective. Before a decision is finalized, state as fact that it failed badly two years on, and ask each participant *independently and in writing* why. Two properties make this a u instrument and not merely a risk exercise: it inverts the frame, so agreement with the group's prior no longer counts as a contribution; and it is written before discussion, so it precedes the cascade.

Read the output as a u measurement: if the independently written failure stories cluster onto one story, u is low on the risk axis. If they scatter across genuinely different mechanisms, u is live. **A pre-mortem that produces one story is not a reassuring pre-mortem. It is a failed measurement of the group.**

Instrument 5 — the shuffled-input control (Chapter 25). The decisive test, and the one most often omitted. Fit the minimal estimation model `Y_{t+1} = a Y_t + b F(U_t) - C_hat_t + e_t` on the real input series; then refit with the `U_t` series randomly permuted in time. If the shuffled series fits as well as the true one, the interaction geometry was not driving the order, and any story told about u in that system is unsupported. **A framework that cannot fail cannot be trusted, and this is the test that lets this one fail.**

11.9 What correlation does to Condorcet

Condorcet's theorem requires independence. Institutions habitually invoke the theorem's conclusion — bigger juries are more accurate — while operating in conditions that violate its premise. The size of the error is computable.

Here is an explicit toy model, defined in this book and not cited from anywhere, so that its assumptions are visible. With probability `rho`, all `N` members receive and act on one common signal (correct with probability `p`), so they vote identically. With probability `1 − rho`, they vote independently, each correct with probability `p`. Majority rules.

`p = 0.60`

rho	N = 3	N = 11	N = 51
0.0	0.6480	0.7535	0.9265
0.2	0.6384	0.7228	0.8612
0.5	0.6240	0.6767	0.7632
0.8	0.6096	0.6307	0.6653
1.0	0.6000	0.6000	0.6000

At `rho = 0`, going from 3 to 51 members buys 28 points of accuracy. At `rho = 0.8` it buys 5.6. At `rho = 1.0` it buys nothing at all: fifty-one perfectly correlated 60-percent voters are a 60-percent voter.

Scale is a multiplier on independence, not a substitute for it, and a large correlated body is not a safer decision-maker than a small one — it is the same decision-maker with a larger travel budget. This is Chapter 8's logarithmic ceiling meeting Chapter 11's parallel wall: raising `N` cannot rescue a system whose `u` is near zero, because `f(u) = ln(1 + (N−1)u)` and `u = 0` kills the `(N−1)` term outright.

11.10 The variance instrument

There is a second, entirely standard effective-count formula from classical statistics, and it is worth having because it is an *independent* instrument — it does not rest on the surplus law's axioms at all, so agreement between the two is meaningful.

Take `N` estimators of the same quantity, each with variance `sigma^2`, with equal pairwise correlation `rho`. The variance of their mean is

`Var(mean) = sigma^2 * [ 1 + (N−1) rho ] / N`

Setting this equal to `sigma^2 / N_eff_var` gives the effective number of independent estimators:

`N_eff_var = N / [ 1 + (N−1) rho ]`

N	rho=0.00	rho=0.10	rho=0.30	rho=0.50	rho=0.80	rho=1.00
3	3.00	2.50	1.88	1.50	1.15	1.00
5	5.00	3.57	2.27	1.67	1.19	1.00
10	10.00	5.26	2.70	1.82	1.22	1.00

```
  20      20.00       6.90       2.99       1.90       1.23       1.00
  50      50.00       8.47       3.18       1.96       1.24       1.00
```

Read the `rho = 0.10` column down. At a mere ten percent error correlation — which most forecasting groups would describe as excellent independence — fifty estimators are worth 8.47. At `rho = 0.30`, fifty are worth 3.18, and the column has nearly stopped responding to `N` at all.

```
CEILING OF THE VARIANCE INSTRUMENT

  As N -> infinity,  N_eff_var -> 1 / rho

  rho = 0.10  ->  ceiling 10 independent estimators, at any headcount
  rho = 0.30  ->  ceiling 3.33
  rho = 0.50  ->  ceiling 2
```

Error correlation imposes a hard ceiling on ensemble accuracy that no amount of hiring can raise. Fifty analysts at `rho = 0.3` and five thousand analysts at `rho = 0.3` are both worth about 3.3 analysts.

Two honesty notes. First, `N_eff_var` and the law's `N_eff = 1 + (N−1)u` are **different quantities from different derivations** — one from second moments of estimator error, one from the surplus axioms. This book does not claim they are equal and does not convert between them. They are used as two independent readings of the same underlying condition, in the way an engineer uses two different gauges on one vessel. Second, the equal-pairwise-correlation assumption is a simplification; real correlation structures are block-structured, and the right treatment there is the eigenvalue spectrum of the correlation matrix. State which you used.

11.11 Maintaining u: driving lambda down

Because u decays, the design question is not "how do we get u?" but "what is our `lambda`, and what does it cost to lower it?"

The devices below are `lambda`-reduction devices. Each attacks a specific cascade mechanism. All are cheap; most are procedural rather than structural, which is why this is usually the least expensive intervention in the whole manual.

```
DEVICE                                  CASCADE MECHANISM IT BLOCKS
-------------------------------------------------------------------------
Independent estimate before             Removes the public history agent 3 would
discussion; form closes before           condition on. Blocks the cascade at its
the meeting opens                        formation.
Silent write-first, then round           Same, at meeting granularity. Every
robin reading of what was written       participant's signal enters the record.
Blind review; second reviewer           Blocks the sequential-observation channel
cannot see the first's score             in review. Cheapest fix in any
until both are submitted                 peer-review system.
Last-speaks-first inversion for          Removes the weighted first signal that
rank                                    forms a cascade on turn one, not turn 3.
Rotating adversarial assignment,        Keeps gamma nonzero without appointing a
protected and voluntary                  permanent opposer (Chapter 10, 10-F).
Rotation across teams and                Attacks Species 2 and 5 at the supply
```

```
lineages; recruitment outside          side: refreshes the input distribution,
the dominant pipeline                   not just the roster.
Minority-report requirement           Forces eta > 0 on the divergence channel;
                                     makes dissent durable, not merely audible.
Decorrelated incentives                 Removes the common-cause term producing
                                         correlated error among differently
                                         trained people.
-------------------------------------------------------------------------
```

The last row deserves emphasis because it is the one that survives all the others. **If every member of a group is paid on the same number, their errors will correlate no matter how differently they were trained, because the incentive is a common signal and a common signal is exactly what a cascade runs on.** Goodhart's Law and Campbell's Law are usually read as warnings about metric gaming. Read here, they are also warnings about u: a single dominant metric is a u-destroyer, because it makes one dimension of the state space matter and flattens the rest.

The cost of maintenance. All of these raise C: a round, a coordination step, drafting time. That cost is real and must be checked against headroom before adoption — Chapter 6's rule stands, never request `r_t > h_t`. What §11.12 shows is that the cost is nearly always small relative to the surplus recovered, *provided the leak is fixed rather than patched once.*

11.12 Why a one-time fix does not work

This is the chapter's final structural result and it is the strongest argument for the maintenance framing.

Take the collapsed institute of §11.7 at `t = 12`, with `Q = 2.3015`, and apply a serious one-time intervention that restores u to 0.42. Compare two futures: one in which a small residual leak of `lambda = 0.03` per quarter remains, and one in which maintenance holds u flat.

`N = 5, eta = 0.55, delta = 0.09, C = 0.15, restored u = 0.42 at t = 12`

	lambda = 0.03		lambda = 0.00 (maintained)	
t	u	Q	u	Q
12	0.4200	2.3015	0.4200	2.3015
16	0.3722	2.8890	0.4200	2.9480
20	0.3298	3.1519	0.4200	3.3910
24	0.2922	3.1980 <- max (3.2021 at t=23)	0.4200	3.6950
28	0.2589	3.1030	0.4200	3.9030
32	0.2294	2.9190	0.4200	4.0460
36	0.2033	2.6800	0.4200	4.1440
40	0.1801	2.4110	0.4200	4.2110
			Q_inf =	4.3578

With a three-percent-per-quarter residual leak, Q peaks at 3.202 at `t = 23` and is back to 2.411 by `t = 40` — within 0.11 of where the intervention started, twenty-eight quarters later. **A leak of three percent per period, which is small enough that nobody would notice it in any single meeting, erases the entire benefit of a serious intervention inside seven years.**

With the leak closed, the same intervention converges to 4.358 — an 89 percent increase over the starting point, and it holds.

The two runs differ in *nothing* except whether the maintenance regime exists. Same restoration, same people, same cost structure, same headroom.

```
DESIGN CONCLUSION

  A one-time diversity intervention on a system with lambda > 0 is a step
  input to a leaky accumulator. It produces a transient and returns.

  The durable quantity is not u. It is 1/lambda -- the machinery that
  keeps u from leaving.
```

That is the chapter's thesis in one line, and it is why pseudo-diversity is correctly classed as a maintenance failure rather than a hiring failure.

The Specification

Requirements are numbered `SPEC 11.n`. Each is checkable.

Group A — measurement

SPEC 11.1 u shall be measured, not assumed. A system that reports a diversity posture without a u estimate is non-conforming.

SPEC 11.2 u shall be re-measured every period, on the same instrument, same item pool, and same pre-discussion ordering. A single measurement is a level; this chapter requires a slope.

SPEC 11.3 `lambda` shall be estimated from two u measurements separated by k maintenance-free periods: `lambda_hat = 1 - (u_1/u_0)^(1/k)`. The maintenance-free condition shall be asserted in writing, with the dates.

SPEC 11.4 The half-life `t_half = ln 2 / (-ln(1 - lambda))` shall be reported alongside `lambda`, in the same units as the reporting period.

SPEC 11.5 `rho_ceil = f(u) / ln N` shall be reported every period. `h_t` alone does not detect a parallel wall (§11.7) and shall not be relied on for that purpose.

SPEC 11.6 At least two independent detection instruments from §11.8 shall be run each period, and at least one of them shall be quantitative.

SPEC 11.7 Where a history of resolvable predictions exists, error correlation `rho` shall be measured and `N_eff_var = N/[1+(N-1)rho]` reported. Where the correlation structure is block-structured rather than exchangeable, this shall be stated and the eigenvalue treatment used instead.

SPEC 11.8 Category blocks shall be declared by evaluative standard (Chapter 10, SPEC 10.1). Blocks declared on any axis other than the one on which judgment varies shall be marked as such, and u measured on both.

SPEC 11.9 Information supply shall be audited alongside the roster: each member's three most view-changing sources of the period, collected independently, with the overlap computed. This detects Species 2, which no roster audit can.

SPEC 11.10 Training lineage, not affiliation, shall be recorded for reviewer, candidate, and appointment pools. Two people at different institutions with one advisor are one lineage.

Group B — protocol

SPEC 11.11 Independent written estimates shall precede discussion on every decision where u is a design input, and the collection form shall close before the discussion opens. This is the chapter's highest-yield requirement and admits no exception on grounds of urgency.

SPEC 11.12 In any sequential review, a later reviewer shall not see an earlier reviewer's score or recommendation before submitting their own.

SPEC 11.13 Where rank is present, the highest-ranked participant shall speak last on substance.

SPEC 11.14 Every substantive decision artifact shall record the strongest case against the decision taken, whether or not anyone advanced it. This is an `eta` requirement, not a u requirement: it makes divergence durable.

SPEC 11.15 Pre-mortems shall be written independently before discussion, and their output read as a u measurement. Convergence of independently written failure stories shall be recorded as a low-u finding, not as confirmation.

SPEC 11.16 No participant shall be required to dissent, and a decline shall be recorded as a decline rather than as a failure (CANON §5.4). Compelled divergence produces low-information objections that raise `C` without raising u.

SPEC 11.17 Where a single composite metric governs the incentives of all members, a u impact assessment shall be recorded. A single dominant metric is a common signal and therefore a cascade substrate.

SPEC 11.18 Adding stages, reviewers, or committee members shall require a prior statement of the expected u contribution of the addition. Additions with no expected u contribution add `C` only (Species 4) and shall be declined.

Group C — reporting and governance

SPEC 11.19 The period report shall contain:

```
PERIOD: _______   N (declared blocks): ____
u: _______   N_eff = 1+(N−1)u: _______   S = f(u): _______
ln N: _______   rho_ceil = S/ln N: _______
lambda: _______   t_half: _______ periods
rho (error correlation): _______   N_eff_var: _______
instruments run this period: ____________________
decisions with recorded dissent: ____/____
dissents that changed the artifact: ____/____
later-stage reversals of earlier stages: ____/____
source-overlap index: _______
```

SPEC 11.20 Postulate 11.1 shall be badged wherever `lambda` is used. It is a modeling assumption of this manual and not a theorem.

SPEC 11.21 Any use of `N_eff` shall badge axiom S3 in the same sentence.

SPEC 11.22 Preregistration (Chapter 25) shall include the shuffled-input control. Its omission shall be treated as a material defect in the analysis.

SPEC 11.23 The report shall state, verbatim: *the surplus functional as deployed is a lower bound on the ensemble's information, not the information itself.*

The Failure Modes

11-A . THE PARALLEL WALL

Clinical. `u` at or near zero across the roster. `N_eff = 1`. Surplus `ln(1) = 0` regardless of headcount.

Damages. `S -> 0`, so the storage law reduces to `Q_{t+1} = (1-delta)Q_t - C_t`: pure decay with a cost draw. No headcount grows it.

Signature. Unanimity as the normal outcome. High confidence, flat information. Meetings whose conclusions are predictable from the agenda.

Measurement. Collision test returning near-zero collision rate; `rho_ceil` below 0.15; error correlation above 0.8.

Fix. Nothing on the `N` lever works (§11.9). Buy `u` directly, per Chapter 9 and Rule 10.4, and install the §11.11 `lambda` devices at the same time — otherwise you have bought a transient (§11.12).

11-B . PSEUDO-DIVERSITY (declared variety, epistemic uniformity)

Clinical. Blocks declared along an axis on which judgment does not vary. The roster is heterogeneous on the declared axis and homogeneous on the deciding one.

Damages. `u` overestimated at declaration and near zero in operation. All downstream planning optimistic.

Signature. A diversity report and a decision record that disagree: many declared categories, no recorded disagreement.

Measurement. Re-declare by evaluative standard (SPEC 11.8) and re-measure. If cross-block `beta^2` does not exceed within-block `beta^2`, the declaration was not epistemic.

Fix. Re-declare the blocks, and treat the resulting number as the real one even when it is far below the previous report.

11-C . THE CASCADE COLLAPSE

Clinical. A group that measured well at formation and measures near zero later, with no personnel change. Individually rational copying (BHW 1992) has removed private signals from the public record.

Damages. u `-> 0` over time; `lambda` large. Because `f` is steepest at zero (§11.0.4), the last increments lost are the most costly.

Signature. Declining variance in stated positions over successive periods. Newer members' positions converging to the group's within one or two cycles. Meetings whose first speaker predicts the outcome.

Measurement. Two u measurements separated by maintenance-free periods. `lambda_hat > 0.10` is a live collapse.

Fix. The §11.11 devices, especially independent-estimate-before-discussion and blind review. These block cascade formation rather than treating its result.

11-D . THE ONE-SHOT INTERVENTION

Clinical. A serious u-raising programme executed once, with no maintenance regime, on a system with `lambda > 0`.

Damages. A transient. `Q` rises, peaks, and returns toward its pre-intervention level (§11.12). The organization concludes the intervention did not work and declines to repeat it, which is the worse outcome.

Signature. A visible improvement 4 to 8 periods after a programme, followed by a slow, unattributed decline.

Measurement. Track u — not `Q` — after the intervention. `Q`'s lag (§11.7) makes it useless as an early indicator.

Fix. Pair every u intervention with a `lambda` intervention, and budget the maintenance before funding the programme. An unmaintained programme is worse than none: it consumes the capital for the next one.

11-E . THE ECHO SUPPLY CHAIN

Clinical. Heterogeneous members, homogeneous inputs. Difference in training does not survive identical evidence.

Damages. u near zero on any question the shared sources address, and normal-looking elsewhere — which makes it intermittent and easy to dismiss.

Signature. Members can articulate different frameworks and cite the same evidence. Positions update simultaneously across the group when one upstream source publishes.

Measurement. SPEC 11.9's source-overlap index.

Fix. Diversify inputs before diversifying the roster. Assign different members to different evidence bases deliberately; require the source list in the decision artifact.

11-F . CORRELATED FORECASTS PRESENTED AS AN ENSEMBLE

Clinical. k estimates with correlated errors, aggregated and reported with ensemble language and ensemble confidence intervals.

Damages. Confidence intervals too narrow by a factor of roughly `sqrt(N / N_eff_var)`. The number reported to decision-makers is wrong in the dangerous direction.

Signature. Ensemble spread much narrower than realized error, repeatedly. Forecasts that miss together.

Measurement. §11.10. Compute `rho` on historical residuals; report `N_eff_var` next to `N` every time.

Fix. Report `N_eff_var`, not `N`, and widen intervals accordingly. Decorrelate at source by varying data, method, and incentive across forecasters.

11-G . THE QUADRATIC COMMITTEE (committee theatre)

Clinical. Additional review stages that add no u. Each stage draws from the same pool and applies the same standard.

Damages. `C` grows with stages while `S` is flat; Brooks' `n(n−1)/2` operates on the cost side only. Chapter 5's crossover arrives early.

Signature. Reversal rate near zero across many decisions. Stages described as "rigorous" that have never changed an outcome.

Measurement. Reversal count per SPEC 11.19. Also: remove a stage in a pilot and measure whether decision quality changes.

Fix. SPEC 11.18. Delete stages that carry no u. Deleting a stage is a `C`-reduction that costs no surplus, which makes it one of the rare unambiguous wins in this manual.

11-H . THE MONOMETRIC INCENTIVE

Clinical. Every member's reward tied to one composite number.

Damages. Correlated error regardless of training diversity. The metric is a common signal, so u decays toward zero on every dimension the metric does not contain. Goodhart and Campbell operating as u-destroyers rather than as gaming risks.

Signature. People with genuinely different backgrounds making identically shaped tradeoffs.

Measurement. Compare u measured on metric-relevant questions against u on metric-irrelevant ones. A large gap is the signature.

Fix. Decorrelate incentives across roles. Accept that this makes performance management harder; that difficulty is the price of the surplus.

11-I . THE MIRROR ROLE

Clinical. A permanent designated contrarian, credited as the group's divergence.

Damages. u credited where none exists (antiparallel is parallel, §11.0.2), and genuine dissent suppressed because the slot appears filled.

Signature. One member's position derivable from the proposal. Shared-ground correlation near −1.

Measurement. Predict the objector's position from the proposal. If reliable, they are a mirror.

Fix. Rotate the role; score by artifact change, not by objection count.

Worked Example

The Standards Review Board

A fifteen-member technical review board at a national standards body, reviewing safety determinations. Five declared specialties. It was constituted eleven years ago and was, at constitution, genuinely diverse: an early u audit conducted in year two returned `u = 0.55`.

Parameters (Chapters 3-5 procedures, quarterly).

```
N = 5        ln N = 1.609438
delta = 0.09     eta = 0.55     C = 0.15     Q at year-2 audit = 4.00
eta ln N = 0.885191           <- the per-period ceiling on eta S
```

Year-2 state.

```
u = 0.55   N_eff = 1 + 4(0.55) = 3.200   (S3 badged: design axiom)
S = ln(3.200) = 1.163151   rho_ceil = 1.163151/1.609438 = 0.7227
eta S = 0.639733    delta Q + C = 0.36 + 0.15 = 0.51
Delta Q = +0.129733   -> growing
Q_inf at this u = (0.639733 - 0.15)/0.09 = 5.441478
```

Healthy. Growing at 0.13 per quarter toward 5.44, with 28 percent of its per-interaction ceiling unused. Nobody had any reason to look further.

What happened next, and why nobody saw it. No one left. No policy changed. Three things happened that no one recorded as events:

1. Scores were collected *in the meeting*, in seniority order, rather than in advance.
2. Second reviewers were routinely given the first reviewer's assessment "for context."
3. All fifteen members' committee stipends were tied to a single throughput metric — determinations closed per quarter.

Items 1 and 2 are cascade substrate (§11.4). Item 3 is a common signal (Failure Mode 11-H). Together they produce a measured `lambda_hat = 0.18` per quarter — half-life 3.49 quarters — which is the decay run computed in §11.7 and reproduced here at the points that matter:

t (quarters)	u	N_eff	S	rho_ceil	Q	
0 (yr 2)	0.5500	3.200	1.1632	0.7227	4.0000	
2	0.3698	2.479	0.9080	0.5642	4.1751	<- Q peak
4	0.2487	1.995	0.6905	0.4290	4.0623	

```
  8           0.1124     1.450    0.3714     0.2308      3.3166
 12 (yr 5)    0.0508     1.203    0.1851     0.1150      2.3015
 16           0.0230     1.092    0.0879     0.0546      1.3240
 20           0.0104     1.042    0.0407     0.0253      0.5110
 23           0.0057     1.023    0.0227     0.0141      0.0248
```

The board's stored order peaked two quarters after its geometry began to collapse, and every report available to its chair during those two quarters said the board was improving.

The audit, run in year 5 (t = 12).

Instrument 1 — collision test. 20 live determinations put to pairs of members separately, in writing. Collisions: 1 of 20. Collision rate 0.05.

Instrument 2 — disagreement audit, last 40 determinations:

```
determinations with any dissent recorded            3 / 40
  dissent from within the deciding specialty        0
  dissent from a different specialty                3
dissents that changed the final artifact            0 / 3
later-stage reversals of an earlier stage           0 / 40
```

Instrument 3 — error decorrelation. The board issues resolvable severity scores; 34 determinations had subsequent field outcomes. Residual correlation across members: `rho = 0.71`.

```
N_eff_var = 15 / [1 + 14(0.71)] = 15 / 10.94 = 1.371
```

Fifteen reviewers worth 1.37 independent reviewers. The surplus-law reading agrees: `u = 0.0508` gives `N_eff = 1.203`. **Two instruments built on entirely different foundations both report a board of fifteen operating as a board of roughly one and a third.**

Instrument 4 — pre-mortem. Written independently on a pending determination. Fifteen failure stories; thirteen of them named the same mechanism. Recorded as a low-u finding per SPEC 11.15.

Instrument 5 — shuffled-input control. The minimal model was fitted on eleven years of quarterly data with Y = the count of standards in force that had been revised on schedule (chosen because it is not a repackaging of U — it contains no measure of disagreement, interaction volume, or composition; Paper E §6.3). The shuffled-input control fitted materially worse than the true series, so the input geometry is doing real work in this system and the analysis is not spurious. Had it fitted equally well, everything above would have been withdrawn.

Diagnosis. Failure Mode 11-C (cascade collapse), driven by 11-H (monometric incentive), presenting as 11-A (parallel wall). Not a hiring failure. The board still contains the diversity it was constituted with; it has stopped being able to express it.

Sustainability check before proposing anything.

```
h_12 = eta ln N − delta Q_12 − C
     = 0.885191 − 0.09(2.3015) − 0.15
     = 0.885191 − 0.207135 − 0.15
     = 0.528056
```

h_12 is large — 0.528 — and rising, because Q has collapsed and `delta Q` with it. **The board has abundant headroom and is dying, which is exactly the combination §11.7 predicted and**

exactly why headroom must never be read alone. The binding constraint is not capacity. It is `rho_ceil = 0.115`: the board is running at eleven percent of its per-interaction ceiling.

Intervention.

```
INTERVENTION SPEC — STANDARDS REVIEW BOARD

A. lambda-reduction (the leak), effective immediately:
   A1. Severity scores submitted to a locked form that closes 24 hours before
       the meeting. No score visible to any member until all are in.
       (SPEC 11.11)
   A2. Second reviewers no longer receive the first reviewer's assessment
       before submitting. (SPEC 11.12)
   A3. Chair and senior members speak last on substance. (SPEC 11.13)
   A4. Stipend decoupled from the throughput metric; throughput retained as a
       reported statistic with no reward attached. (SPEC 11.17)
   A5. Minority-report field mandatory on every determination artifact.
       (SPEC 11.14)  Target: lambda <= 0.03.

B. u-restoration (the level), quarters 1-4:
   B1. Rotating within-specialty adversarial assignment, protected and
       voluntary; declining is recorded as a decline. (SPEC 11.16)
   B2. Two members per quarter seconded to review outside their specialty.
   B3. Source-overlap index measured and inputs deliberately varied.
       (SPEC 11.9)  Target: u >= 0.42.

C. Cost control:
   C1. C monitored monthly. Estimated increase +0.04 per quarter.
       At C = 0.19, h = 0.885191 - 0.207135 - 0.19 = 0.488056. Ample.
   C2. One review stage deleted: the pre-circulation summary stage, which had
       produced zero reversals in 40 determinations (SPEC 11.18). Estimated
       C saving -0.02, offsetting half the increase.
```

Predicted trajectory. With u restored to 0.42 and lambda held at 0.03:

```
S at u=0.42 = ln(1 + 4(0.42)) = ln(2.68) = 0.985817
eta S = 0.542199   rho_ceil = 0.6125
Delta Q at t=12 = 0.542199 - 0.207135 - 0.15 = +0.185064

 t          u        Q  (lambda = 0.03)      Q  (lambda = 0, maintained)
-------------------------------------------------------------------------
12       0.4200        2.3015                      2.3015
16       0.3722        2.8890                      2.9480
20       0.3298        3.1519                      3.3910
24       0.2922        3.1980  (max at t=23)       3.6950
32       0.2294        2.9190                      4.0460
40       0.1801        2.4110                      4.2110
                                                Q_inf = 4.3578
```

This is the decision the board is actually making. Both columns start from the same interven-

tion. The left column leaks three percent per quarter and is back near its starting point after seven years. The right column holds and converges to 4.358, an 89 percent gain.

The difference between them is not effort, budget, or personnel. It is whether items A1 through A5 are still being performed in year eight. **The intervention that matters is not the one that raises u; it is the one that keeps u from leaving, and it is the one most likely to be quietly dropped, because it is procedural, invisible when working, and mildly annoying to everyone every single week.** Star and Ruhleder's law applies to protocols as much as to plumbing.

Falsification criteria, preregistered and frozen (Chapter 25 form).

```
PREDICTIONS, at quarter 4 of the intervention:
  P1. measured u >= 0.30
  P2. measured lambda <= 0.06
  P3. error correlation rho <= 0.45, so N_eff_var >= 2.1
  P4. determinations with recorded, artifact-changing dissent >= 6 / 40

FALSIFIED IF:
  (a) u < 0.20 despite full protocol compliance, OR
  (b) lambda > 0.12 despite full compliance, OR
  (c) Y has not improved against the frozen persistence baseline, OR
  (d) the shuffled-input control fits as well as the true input series.

REQUIRED BASELINES: persistence; direct-input; matched autoregressive;
                    shuffled-input control.

ORDER PROXY Y: count of standards in force revised on schedule. Justified
as NOT a repackaging of U -- it contains no measure of disagreement,
interaction volume, or board composition. (Paper E, §6.3.)

COMPLIANCE IS A MEASURED VARIABLE, NOT AN ASSUMPTION. Protocol adherence
is logged per determination. A failed prediction under partial compliance
falsifies nothing and shall be reported as an implementation failure, not
as evidence against the model.
```

What this example does not claim. It does not claim the board was badly constituted, badly staffed, or badly intentioned; the year-2 audit says otherwise on all three. It does not claim the surplus law explains the board's performance — that is what the shuffled-input control is there to test, and it could have failed. It claims one thing, narrowly and with arithmetic:

a group that was genuinely diverse became arithmetically indistinguishable from a single reviewer, in five years, through three procedural choices that nobody recorded as decisions — and the only reason it was caught is that someone measured a quantity most institutions have never measured at all.

Chapter 12 leaves the roster and moves to the room. The reason follows directly from this chapter: several of the strongest `lambda`-reduction devices — speaking order, sightline reciprocity, whether the quiet voice can be heard at all — are not policies. They are properties of furniture and acoustics, and they enforce themselves continuously without anyone administering them.

Part III — Engineering Space

Chapter 12 . The Room as Circuit Element

Parts I and II were about quantities you could only reach through people: how fast stored order leaks, how much of an encounter survives it, how different the participants actually are. Every one of those quantities has a physical substrate, and the substrate is usually a room.

This is the first physical chapter. From here to Chapter 17 the manual works in metres, decibels, lux, and degrees Celsius, and it does so for one reason: **the coefficients of the storage law are not fixed properties of a group; they are partly properties of the place the group is standing in, and place is the cheapest thing in the system to change.**

You cannot legislate a person into being more different from their colleague. You can move a wall. The wall is nine hundred kilograms of gypsum and it enforces its policy continuously, without a budget line, without a meeting, and without anyone's permission being renewed each quarter. That is the argument of Part III in one sentence.

The Law

What is proven

The storage law is a theorem about a recursion, and it is stated in CANON §1.2 exactly as Paper E proves it:

```
    Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t        S_t = F(U_t) = ln(1 + (N−
1) U_t)

    0 <= S_t <= ln N                                  (Paper D, Corollary 2.2)
     h_t = eta ln N − delta Q_t − C_t                        (Paper E, §6.4 −−
the headroom gauge)
```

Nothing in this chapter alters that. Nothing in this chapter is permitted to alter that. What this chapter does is supply an *interpretation of the coefficients for a bounded physical volume containing people*, and interpretation is not proof.

What is postulated

POSTULATE 12.1 (badged as a postulate, not a theorem): a room is a circuit element. It has its own delta, eta, N, u and C, and those values are measurable by direct observation. This is a modelling assumption in exactly the sense CANON §1.4 describes: the claim is about *model form*, not about physics. The storage law is scale-free in form; asserting that a 40-square-metre volume is one of the scales at which the form applies is a choice, and it must be earned by measurement in each case, not assumed.

POSTULATE 12.2 (badged): optional activity is the observable proxy for eta. Gehl's taxonomy divides what happens in a space into activities that would happen anyway and activities that happen only when conditions permit. The second category is behaviour that the space *caused*. Since eta is defined as the fraction of interaction that converts into durable structure, and since a space that causes voluntary lingering is a space where interaction has somewhere to go, we take

the optional-activity rate as eta's field instrument. **This is a proxy, and a proxy is a bet.** §12.3 states precisely what would falsify it.

POSTULATE 12.3 (badged): a room holds a small but non-zero Q of its own. Furniture arrangements that persist, worn paths, the shelf everyone knows to check, the corner that is understood to be for quiet work — these are stored order held in the physical medium rather than in anyone's head. They decay (delta) and they are replenished (eta S). Treating them as a Q term is a modelling choice. It is a productive one, because it explains why rooms that are "reset" nightly by cleaning staff never accumulate any convention.

The honesty clauses, restated because this chapter leans on them

- **As deployed, the surplus functional is a LOWER BOUND** on the ensemble's information, not the information itself. When §12.10 computes a room's N_eff, it is computing a floor.
- **Paper A is pure mathematics and makes no empirical or metaphysical claim.** A room does not "have surplus". A group of people interacting in a room may be *modelled* as having it.
- **S3, the affine effective count, is the design axiom** — it is chosen, not derived. Every N_eff figure in this chapter inherits that choice.
- **The triangle inequality fails for geometric triples.** Do not chain room-to-room surplus estimates as if they added along a path.
- Numbers in this chapter that come from the observational literature (Whyte, Gehl, Hall) are **borrowed and nominal**. They were measured in specific cultures, in specific decades, with specific populations. They are starting points for your own measurement, not constants.

The load-bearing statement of the chapter: a room does not create interaction, it sets the coefficients under which interaction either converts or dissipates.

The Engineering

12.1 The four terms, located in a physical volume

Take the storage law and ask, for a specific room, what each symbol is pointing at.

Term	In a room	Typical field proxy	Units of the proxy
`Q_t`	Durable relation, standing agreement, learned convention, the arrangement that persists	Count of standing cross-group working agreements; count of persistent furniture states	dimensionless count
`S_t`	Interaction surplus generated by encounters in the volume this period	`ln(1 + (N−1)U_t)` computed from an encounter log	nats

Term	In a room	Typical field proxy	Units of the proxy
`eta`	Rate at which an encounter becomes durable relation or decision	attributable `dQ` per unit `S`	dimensionless
`delta`	Rate at which the room's stored order decays without maintenance	`1 − Q_1/Q_0` measured under zero input	per period
`C_t`	Friction of using the volume at all	booking overhead + travel + setup + acoustic fatigue + interruption	nats-equivalent, see §12.8
`N`	Genuinely distinct activity categories the room supports	category audit, §12.10	count
`u`	Non-parallelism of the people the room actually brings into contact	disagreement rate / error decorrelation across observed pairs	in `[0,1]`

Two of these deserve immediate comment.

`N` for a room is not the occupancy. It is the number of genuinely distinct *kinds* of use the volume supports. A room with twelve identical desks in twelve identical postures under one lighting condition offers one category, not twelve. Chapter 7's arithmetic applies without modification: `N_eff = 1 + (N−1)u`, and at `u = 0`, `N_eff = 1`. **Ten parallel windows are one window; twelve parallel workstations are one workstation.**

`C` for a room is the term nobody prices. It includes the six minutes to find a free room, the four minutes of cable fumbling, the eleven-minute walk, the cognitive cost of raising your voice over a reverberant field for fifty minutes (Chapter 14 prices that one exactly), and the interruption cost paid by everyone the room's traffic passes. §12.8 gives a tally sheet.

12.2 Whyte's method: direct observation as instrumentation

William H. Whyte's *The Social Life of Small Urban Spaces* (1980) is the methodological ancestor of this chapter, and its method matters more than any of its findings. Whyte and the Street Life Project did something that had not been done systematically: they pointed cameras at plazas and counted what people did.

They did not survey. They did not interview. They did not model. **They watched, they counted, and they timed, and the findings that came out contradicted a great deal of what designers believed they knew.**

The method decomposes into four instruments, and all four transfer directly to interior space:

Instrument 1 — Time-lapse. A fixed camera at a fixed interval (Whyte used film; a phone on a tripod at one frame per five or ten seconds is now trivially cheap). Time-lapse converts occupancy

into a legible signal. Standing patterns that are invisible in real time — that nobody ever sits on the north bench, that the queue always kinks left, that the room empties at 15:40 not 16:00 — appear immediately.

Time-lapse is also the honest instrument, because it records what you did not think to look for. A count sheet records only your hypotheses.

Instrument 2 — The count. At fixed times, on a fixed plan, mark every person's position and what they are doing. Whyte's core counts were: total people, people in groups (as opposed to alone), and the sex ratio. This manual adds Gehl's activity category (§12.3).

Instrument 3 — The trace. Wear patterns, desire lines, displaced furniture, the chair that is always at 40 degrees to the table, the ledge with the polished patch. Traces are integrated occupancy — they are the room's own Q, written into its surface. Book One, Node 02 taught the operator to read a desire line as a measured gradient. Here the engineer reads it as a time-integral of use that no observer had to be present for.

Instrument 4 — The intercept, used last and sparingly. Ask people why. Whyte's own practice was to observe first and ask later, because asking first contaminates the observation and because people are poor reporters of their own spatial behaviour. Nobody says "I sat here because it is seventeen inches high and I can see the door." They say "it seemed nice."

Whyte's findings that transfer to interiors State these as borrowed and nominal. They were measured on New York plazas in the 1970s.

- **"People sit where there are places to sit."** Whyte's own summary of his largest finding, and it is deliberately flat. The amount of *sittable space* predicted plaza use far better than any aesthetic or architectural variable he tested.
- **Movable seating outperforms fixed seating**, even when the mover barely moves it. The amenity is not the position; it is the *demonstrated ability to choose the position*. This is a direct antecedent of Chapter 16's rule that adjustability beats optimisation.
- **Triangulation**: some external stimulus — a performer, an unusual object, a view, a shared problem — that gives strangers a legitimate reason to speak to each other. §12.6.
- **Self-congestion**: people stop and talk *in the traffic*, at the mouth of the space, not in the calm interior. The most-used part of a plaza is the part nearest the street.
- **Sun, wind, trees, water, food.** Each independently associated with use.
- **The proportion of women as a quality index.** Whyte reported that where people had a genuine choice of spaces, women were more selective, so a high proportion of women in a space indicated a space that was working. Flag this hard: it is a 1970s observation about a particular city, and it is a correlate, not a mechanism. It earns a place here only because Woolley et al. (2010) independently found proportion of women among the correlates of the collective intelligence factor c (Chapter 18). **Two unrelated literatures pointing at the same index is worth a measurement; it is not worth a belief.**

The sittable-space ratio The plaza standards adopted in New York after Whyte's work specified, in imperial units, roughly one linear foot of sitting space per thirty square feet of open space. Converted:

```
1 ft of seating per 30 ft^2
  = 0.3048 m per 2.7871 m^2
```

```
= 0.1094 m of seating per m^2
-> 1 metre of seating per 9.14 m^2 of open space
```

Open floor area	Minimum sittable edge (Whyte ratio)
40 m^2	4.37 m
100 m^2	10.94 m
200 m^2	21.87 m
500 m^2	54.68 m
1000 m^2	109.36 m

Use this as a *floor*, not a target, and use it knowing its provenance: it is a planning regulation derived from outdoor observation in one city. It is in this manual because it is one of the very few quantitative space standards that was derived from counting people rather than from committee. Chapter 13 will show a geometric trick for meeting the ratio without spending floor area, and Chapter 15 will show why the seating wants to be on the perimeter.

Whyte's dimensional recommendations for the seating itself — nominal, verify locally:

Element	Whyte's nominal	Metric	Note
Seat height	~17 in	~430 mm	The usable band is wide; ~400-500 mm is sittable
Ledge depth, one-sided	>= 16 in	>= 400 mm	Below this, people perch, they do not sit
Ledge depth, two-sided	>= 30 in	>= 760 mm	Two backs to back without contact
Movable chairs	as many as possible	–	The single highest-yield item in his findings

12.3 Gehl's taxonomy and the eta proxy

Jan Gehl's *Life Between Buildings* (1971) divides outdoor activity into three classes. The division is the most useful single idea in this chapter.

Class	Definition	Dependence on environmental quality
Necessary	Activities people must do: go to work, wait for the bus, walk to the printer, attend the mandatory meeting	Occurs almost regardless of quality
Optional	Activities people do only if they want to and conditions invite it: sit, stand and look, linger, take the long way	Occurs **only** when conditions are favourable

Class	Definition	Dependence on environmental quality
Social	Activities that depend on the presence of others: conversation, greeting, children playing, simply watching	Largely *resultant* — they arise out of the other two

Gehl's central observation is that the third class is derived. Social life is not something you can install. It is what happens when necessary and optional activity coexist in the same volume for long enough. **"Something happens because something happens."**

For this manual the operational consequence is precise:

```
OAR := optional_count / (necessary_count + optional_count)

OAR is the field proxy for eta.
```

The logic: necessary activity is S that would have occurred anyway; the room gets no credit for it. Optional activity is behaviour the room *caused*. Social activity is optional and necessary activity that has begun to interact — which is precisely the raw material eta converts. A room where OAR is 0.10 is a corridor with chairs in it. A room where OAR is 0.35 is generating its own traffic.

What would falsify Postulate 12.2. If a room's OAR rises and its attributable dQ per unit S does not, the proxy is broken for that room. That is a real and common case: a beautiful lounge in which people linger pleasantly and nothing is ever decided, recorded, or carried out has high OAR and eta near zero. Chapter 4 named this the Leaky Vessel. **OAR measures the opportunity for conversion, not the conversion.** Measure both, and if they diverge, believe the dQ.

Gehl's distance thresholds — borrowed, nominal, and worth measuring yourself Gehl reports approximate distances at which the human sensory apparatus does different work. Treat as order-of-magnitude:

Distance	What is available
~100 m	A person is identifiable as a person; gait and body language readable — the "social field of vision"
~50-70 m	Sex, age, general activity readable
~25 m	Facial expression and emotional state readable
~7 m	Full sensory contact; detail, small gesture, ordinary speech comfortable
~1-3 m	Conversation at normal effort with full non-verbal channel

This ladder is the reason large open-plan floors do not produce more contact than small ones. Beyond roughly 25 m, other people are scenery. Beyond roughly 7 m, they are not conversational partners without a deliberate approach, and a deliberate approach has a social cost that most

people will not pay for a marginal exchange. **A 900 m^2 open floor does not contain 900 m^2 of social opportunity; it contains the sum of its sub-7-metre neighbourhoods, and if it has no sub-structure it contains very few of them.** Chapter 13 is entirely about how to install that sub-structure without adding walls.

12.4 Hall's proxemics: the geometry of u at conversational range

Edward T. Hall's *The Hidden Dimension* (1966) established that interpersonal distance is structured, not continuous, and that the structure carries meaning. His four zones, each with a close and a far phase, for North American middle-class adults:

Zone	Range	Character	What the channel carries
Intimate	0 - 0.45 m	Touch, heat, scent detectable	Involuntary signal; intrusion if unconsented
Personal	0.45 - 1.2 m	Arm's length; detail of expression	Confiding, close collaboration
Social	1.2 - 3.6 m	Business range; whole body visible	Working conversation, group of 3-8
Public	3.6 m +	Formal; voice must be raised	Address, performance, one-to-many

Hall himself insisted these are culturally specific. The zone boundaries differ substantially across cultures and the differences are large enough that mismatched norms read as aggression or coldness. Do not ship this table as a constant. Ship it as a hypothesis to check against the people who will actually use the room.

Why proxemics is a u-instrument Chapter 9 established that u is the parameter you can move fastest, and that systems destroy u without noticing. Proxemics is where architecture touches u directly, because **the distance structure of a room determines which pairs of people are ever in a channel that can carry disagreement.**

Consider a room's occupants as a set of pairs. For each pair, the room's geometry assigns a typical separation, and Hall's table assigns that separation a channel type:

- **Public-range pairs (> 3.6 m)** exchange announcements. Announcement is a one-way channel. It cannot carry the hesitation, the partial objection, or the "hmm" that is the raw form of most genuine difference. Pairs at public range contribute close to zero to measured u because their disagreements are never expressed and therefore never counted.
- **Social-range pairs (1.2 - 3.6 m)** are where working difference lives. This is the productive band and almost everything in Chapters 14 and 17 is aimed at maximising the number of pairs in it.
- **Personal-range pairs (0.45 - 1.2 m)** carry the most bandwidth but cannot be sustained across a group; putting eight people at personal range produces a huddle that excludes anyone outside it, and huddles are u-destroyers because they select for the already-agreeing.

- **Intimate-range imposition (< 0.45 m without consent)** raises C sharply and immediately. A room that forces unconsented intimate proximity — a crowded lift lobby used as a meeting point — produces avoidance, and avoidance is a C term that never appears on any budget.

Design rule, stated as a rule and not as a metaphor: count the pairs your geometry places in the social band. That count, not the headcount, is the room's conversational capacity.

For a round table of diameter `D` seating `n`, the maximum separation is `D` (across the diameter) and the seat pitch is `pi*D/n`. Working the arithmetic at a 600 mm minimum pitch:

n	Minimum D for 600 mm pitch	Chosen D	Actual pitch	Max separation	All pairs in social band?
4	0.76 m	1.6 m	1.26 m	1.60 m	yes
6	1.15 m	1.6 m	0.84 m	1.60 m	yes
8	1.53 m	1.6 m	0.63 m	1.60 m	yes, but pitch is tight
10	1.91 m	2.0 m	0.63 m	2.00 m	yes
12	2.29 m	2.4 m	0.63 m	2.40 m	yes
13	2.48 m	2.8 m	0.68 m	2.80 m	yes
16	3.06 m	3.2 m	0.63 m	3.20 m	at the edge (3.2 of 3.6 m)

A round table stops being a single conversation at about sixteen people, and it stops for a geometric reason: the diameter crosses out of Hall's social band. Chapter 17 takes this further. Chapter 14 will show that the acoustic constraint bites before the proxemic one.

12.5 Sittable space, and why it is the highest-yield variable

Whyte's finding is worth restating in the language of the law. Seating does three things simultaneously:

1. **It converts throughput into dwell time.** A person walking through a room at 1.3 m/s is present for four seconds. A person seated is present for twenty minutes. S is generated per unit of co-presence, so seating multiplies S by two to three orders of magnitude per person admitted to the room.
2. **It raises u by admitting a wider population.** Standing-only spaces select for the physically comfortable standing, the unhurried, and the socially confident. Seating admits people who would otherwise not stay, and the people who would otherwise not stay are — by construction — a different population from those who would. **Accessibility is a u-instrument before it is anything else.**
3. **It lowers C.** Standing in a conversation you would like to leave costs more than sitting in one, because standing has no graceful exit protocol other than leaving. Seating with an easy exit (aisle access, no bag-shuffle, no one to climb over) reduces the cost of *entering* a conversation, because the cost of entering is dominated by the anticipated cost of leaving.

Point 3 is under-appreciated and worth a rule of its own: **people join conversations whose exit cost they can estimate.** An alcove with one entrance is a trap; an alcove with a clear line out is an invitation. Chapter 15 develops this as the threshold principle.

Seating typology, ranked by observed yield in Whyte's work and by the reasoning above:

Type	Dwell	Choice	Exit cost	Notes
Movable chair	high	high	low	Highest yield. The move need not be large
Movable chair + movable small table	high	high	low	Enables the object-mediated conversation
Ledge / step, 400+ mm deep	medium	high	very low	Cheapest sittable metre; no maintenance
Bench, backless	medium	medium	low	Sittable from both sides at 760+ mm
Bench, with back	high	low	low	Comfortable, but fixes orientation and so fixes who talks to whom
Fixed banquette in an alcove	high	low	medium	Excellent for two to four; poor if the only option
Fixed rows facing one way	low (optional), high (necessary)	none	high	An auditorium. Optional activity near zero by design

The last row is the important one. **Fixed unidirectional seating sets the room's u toward zero by geometry: it establishes one front, and a room with one front has one speaker and n-1 audiences.** Chapter 17 prices this exactly.

12.6 Triangulation

Whyte's term for the mechanism by which an external stimulus links strangers. A street musician, a peculiar sculpture, a view of something happening, a stuck vending machine — anything that gives two people who have no relationship a legitimate, low-cost, deniable reason to speak.

Triangulation matters here because it is a **u-preserving contact mechanism.** Most engineered contact selects on similarity: the interest group, the affinity network, the project team. Contact selected on similarity buys contact at u near zero, which by Chapter 7's arithmetic is worth `ln(1) = 0`. Triangulation selects on *co-presence at a stimulus*, which is orthogonal to similarity. It is one of the few cheap devices that produces contact between people who have nothing else in common.

Triangulating objects, ranked by robustness:

1. **A shared, visible, changing artefact.** A wall-mounted display of live system state, a physical kanban, a whiteboard nobody erases, a plot on the wall. Also a stigmergic medium in

the sense of Chapter 21 — it does double duty.

2. **A shared consumable with a queue.** Coffee, water, food. The queue is not a defect; the queue is dwell time at conversational range with a legitimate reason to be silent or to speak. **A coffee machine with a thirty-second cycle is a worse social instrument than one with a ninety-second cycle**, and this is the only place in this manual where a slower machine is specified deliberately.
3. **A view.** Ulrich (*Science*, 1984) is the strongest single piece of evidence in this literature that a window view affects a hard outcome — post-surgical recovery — and Chapter 16 handles it. As a triangulator, a view of *activity* outperforms a view of scenery, because activity changes and change gives people something to remark on.
4. **A problem.** The most reliable triangulator ever built is a shared obstacle. Use with care: manufacturing obstacles to create bonding is a well-known way to raise C while pretending to raise eta.

12.7 Sun, food, water, trees

Whyte found each of these independently associated with use. The mechanisms are partly thermal and physiological (Chapter 16), partly attentional. Stated as room-scale requirements:

- **Sun / daylight.** People move to follow it. In an interior, daylight from more than one direction and at more than one intensity gives occupants a real choice, and choice is the amenity. A room with one uniform illuminance offers `N = 1` in the lighting category.
- **Food.** The single most reliable optional-activity generator available. It supplies a legitimate reason to be present, a reason to be still, a reason to be slow, and a shared object.
- **Water.** Whyte specifically noted that people touch water, sit near it, and that its sound masks conversation. The masking is a real acoustic function — see Chapter 14 on background level — and is the one case where deliberately raising background noise is correct: a low, broadband, steady source raises privacy without raising annoyance the way an HVAC tone does.
- **Trees / planting.** Beyond Ulrich's finding, planting does three engineering jobs: it provides sub-division without enclosure (Chapter 13), it absorbs and scatters high-frequency sound slightly (do not overstate this — the absorption coefficient of foliage is real but small), and it gives a visual focus that permits a person to be present without appearing to wait for someone.

12.8 The room's C term: a tally sheet

`C_t` is the term most often left blank, and it is the term that determines whether the room's headroom is positive. Chapter 5 established the general form; here is the room-specific tally.

```
C_room = C_access + C_setup + C_effort + C_interrupt + C_exclusion + C_maint
```

Component	What it is	How to measure it
`C_access`	Time to reach and to secure the room	Median door-to-door minutes x participants; booking-failure rate
`C_setup`	Time from arrival to first productive minute	Stopwatch, ten consecutive sessions

Component	What it is	How to measure it
`C_effort`	Physiological cost of using the room: vocal effort, listening effort, thermal discomfort, glare	Chapter 14 §14.6 gives the acoustic component in dB of required vocal effort; note reported hoarseness
`C_interrupt`	Cost imposed on people *outside* the room by its traffic and noise	Count of interruptions per hour at the nearest workstations
`C_exclusion`	Value of the contributions the room's geometry or acoustics silenced	The `ln N` loss computed in Chapter 14 §14.7
`C_maint`	Cleaning, resetting, repairing, replacing	Hours per week

The rule from Corollary 3.4 applies unchanged: if `C_t > eta ln N − delta Q_t`, no arrangement of the people in the room prevents order from decreasing. A room can be so expensive to use that no seating plan, facilitation technique, or agenda saves it. When you find such a room, the correct move is not to redesign the meeting. It is to cut `C` or to stop using the room.

Two `C` components are chronically invisible:

`C_interrupt` is paid by people who are not in the room and who therefore never appear in the room's evaluation. This is the classic externality structure and it produces the classic outcome: the room is over-used because its full cost is not borne by its users. Chapter 21's Ostrom principles are the correct treatment — the boundary of the resource must include everyone who bears its cost.

`C_exclusion` is paid by the people whose contribution the room suppressed, and they cannot report it, because a contribution that was never made leaves no trace. This is the deepest measurement problem in Part III. Chapter 14 solves it for the acoustic case by computing the loss from physics rather than from self-report. That is why acoustics is the chapter that follows this one.

12.9 The room's delta

`delta` for a room is the rate at which its stored order decays without maintenance. Chapter 3's estimation method applies directly: measure `Q` at two times under zero input, and

```
delta = 1 − (Q_1 / Q_0)          t_half = ln 2 / (−ln(1 − delta))
```

Room-scale `Q` and its decay:

Stored order	Decays because	Nominal delta driver
Furniture arrangement that works	Cleaning reset, the next booking, a single inconsiderate user	Very high without a rule; near zero with a fixed floor marking

Stored order	Decays because	Nominal delta driver
Whiteboard content	Erasure	Total on erasure; delta is effectively a step function
Convention ("the corner is for quiet")	Occupant turnover, absence of a visible marker	Tracks personnel turnover
Wear paths, patina	Resurfacing, refurbishment	Very low; this is the most durable Q a room holds
Artefacts left in place (the shared book, the tool board)	Removal, tidying, theft	Depends entirely on whether removal is anyone's job

The single highest-leverage delta intervention in a room is a floor marking. A taped or inlaid outline showing where the table and chairs belong converts a convention held in fragile human memory (high delta, resets on turnover) into a convention held in the physical medium (very low delta, survives everyone). It costs almost nothing. It is Chapter 21's rule — put the state in the medium, not in a coordinator — applied at the smallest possible scale.

The counter-case from Chapter 3 applies here too: a room with no forgetting cannot learn. A room whose arrangement is welded in place holds its Q perfectly and can never be improved by its occupants. The correct target is not delta = 0; it is delta low enough to survive a week of disuse and high enough that a better arrangement can displace the current one within a month. In practice: **the furniture should be movable, and its default position should be marked.** That is low delta on the default and non-zero delta on the deviation, which is exactly the asymmetry you want.

12.10 The room's N and u

To compute N for a room, audit categories of use, not counts of anything.

Procedure — the category audit.

1. List every distinct *posture-plus-purpose-plus-condition* the room currently supports. A posture is standing, perching, sitting upright, sitting reclined, lying. A purpose is solo focus, pair work, small-group deliberation, presentation, waiting, eating, transit. A condition is bright/dim, quiet/lively, enclosed/open, warm/cool.
2. Strike any entry that is not actually available — a category that requires moving 90 kg of furniture is not on the menu.
3. Strike any entry that duplicates another in every respect that matters to a user. **Two identical four-person tables are one category.**
4. The surviving count is N.

Most rooms audit at `N = 1` or `N = 2` and their owners are surprised. A standard meeting room offers: sitting upright, at a rectangular table, under uniform light, for group talk. That is one category. Its `N_eff` is therefore at most 1 regardless of who is in it, and by Corollary 2.2 its per-interaction surplus ceiling is `ln 1 = 0` in the use-category dimension. The room contributes nothing to S on its own account; all surplus must come from the people, and the room's job reduces to not destroying it.

To estimate u for a room, you are estimating the non-parallelism of the people the room actually brings into contact — not of the organisation, but of the pairs the geometry produces.

Procedure — the contact-pair audit.

1. From the time-lapse or count sheets, list every observed conversational pair over the survey period, with duration.
2. For each pair, score whether the two are in the same team, the same discipline, the same training lineage, the same reporting line. Chapter 9's criteria apply.
3. Weight by duration. Compute the fraction of conversational minutes spent between pairs that differ on at least two of the four axes. Call it `p_cross`.
4. Use `p_cross` as the room's first-order u proxy, and state that you have done so.

This is a crude proxy and must be badged as one. It substitutes categorical difference for geometric non-parallelism, which is exactly the substitution Chapter 11 warns against — demographic or organisational variety is not epistemic variety. Use it as a screening instrument only, and follow with the disagreement audit from Chapter 11 §11.4 before believing any number it produces.

Worked instance of the arithmetic, using Chapter 7's formula and the ceiling from Paper D:

u	`N_eff = 1 + (N−1)u` for N=8	`F = ln N_eff`	Fraction of ceiling `ln 8 = 2.0794`
0.00	1.00	0.0000	0.0%
0.10	1.70	0.5306	25.5%
0.25	2.75	1.0116	48.6%
0.50	4.50	1.5041	72.3%
0.75	6.25	1.8326	88.1%
1.00	8.00	2.0794	100.0%

Read the concavity: **moving u from 0 to 0.25 buys 48.6% of the entire available ceiling; moving it from 0.75 to 1.00 buys 11.9%.** The first genuinely different contact the room produces is worth more than the last four. A room that produces *any* cross-lineage contact is doing most of the work a room can do.

12.11 Occupancy, Brooks, and the sizing of a room

Chapter 5 established that coordination cost grows with the number of channels, `n(n−1)/2`, while per-interaction surplus is capped at `ln n`. For a room, this is a sizing law.

n	channels `n(n−1)/2`	`ln n`	channels per nat
2	1	0.693	1.44
3	3	1.099	2.73
4	6	1.386	4.33
5	10	1.609	6.21
6	15	1.792	8.37
8	28	2.079	13.47
10	45	2.303	19.54
12	66	2.485	26.56
13	78	2.565	30.41

n	channels n(n−1)/2	ln n	channels per nat
16	120	2.773	43.28
20	190	2.996	63.42
25	300	3.219	93.20

Cost grows quadratically while surplus grows logarithmically, so every room has a size beyond which adding a person subtracts. The table does not tell you where that point is — that depends on `eta` and on the per-channel cost in your setting — but it tells you the shape, and the shape is brutal. Going from 8 to 16 occupants buys 33% more surplus ceiling and costs 4.3 times more channels.

The room's job is to *decompose* n rather than to hold it. A twenty-person room that is acoustically and geometrically one space carries 190 channels. The same twenty people in a room sub-divided into four alcoves of five carry `4 x 10 = 40` internal channels plus whatever crosses between alcoves. **Sub-division is not a compromise on connectivity; it is the only way a large room stays affordable.** Chapter 13 is the geometry of doing that without walls.

12.12 What a room cannot do

Three limits, stated plainly, because Part III is where practitioners over-claim.

A room cannot manufacture `u`. If everyone admitted holds the same training, the same incentives, and the same information, the geometry is irrelevant: `u = 0`, `N_eff = 1`, `F = 0`. Chapters 9 and 11 own that problem and no amount of furniture solves it.

A room cannot raise `eta` above what the surrounding practice permits. A perfect room feeding an organisation with no decision log, no artefact, and no memory produces beautiful conversations that vanish. The room raises `eta` from what it would have been; it does not set the ceiling.

A room cannot make a negative-headroom system positive on its own. The Worked Example at the end of Chapter 14 demonstrates exactly this: full acoustic treatment closes 75% of the gap to `h = 0` and the last 25% has to come from shedding `Q`. **Physical intervention is necessary and frequently insufficient, and a manual that pretended otherwise would be selling furniture.**

The Specification

Numbered, checkable, buildable. `R12.n` items are requirements; `M12.n` items are measurements. Every requirement names the term of the law it acts on.

Requirements

R12.1 (acts on `S`, `u`). Sittable space shall be provided at not less than 0.109 linear metres per square metre of open floor area (the Whyte ratio). For a 120 m^2 room this is 13.1 m.

R12.2 (acts on `u`, `C`). Not less than 40% of seating shall be movable by one person without assistance, defined as under 8 kg and free of cabling.

R12.3 (acts on `S`). Seat height shall lie in 400-500 mm; sittable ledge depth shall be not less than 400 mm one-sided and not less than 760 mm where sitting from both sides is intended.

R12.4 (acts on `N`). The room shall support not fewer than three distinct posture-plus-purpose-plus-condition categories, verified by the category audit of §12.10 and recorded on the worksheet.

R12.5 (acts on `u`, `C`). For every seat, the maximum separation to any other seat in the same intended conversational group shall not exceed 3.6 m (Hall's social band). Groups requiring greater separation shall be treated as separate groups and given separate sub-volumes.

R12.6 (acts on `C`). Every seat shall have an exit path not requiring another occupant to move.

R12.7 (acts on `eta`). The room shall contain at least one triangulating element per 40 m^2: a shared changing artefact, a consumable with a queue, a view of activity, or a shared physical tool. Static decoration does not count.

R12.8 (acts on `eta`). The room shall contain at least one writable, non-erased surface of not less than 2 m^2, and a defined owner responsible for photographing and archiving it before any erasure. **A whiteboard with no archive protocol has eta = 0 by construction.**

R12.9 (acts on `delta`). The default position of every movable element shall be marked on the floor or the surface it stands on. Markings shall be inlaid or taped, not printed on paper.

R12.10 (acts on `delta`). Room reset shall be specified as a written, posted, 60-second procedure, and it shall be the responsibility of the last occupant, not of cleaning staff.

R12.11 (acts on `C`). Median door-to-door access time for the intended user population shall not exceed 3 minutes. Where it does, the room's `C_access` shall be tallied explicitly and subtracted in the headroom calculation before the room is approved.

R12.12 (acts on `C`). Setup time from arrival to first productive minute shall not exceed 90 seconds, measured over ten consecutive real sessions, not in commissioning.

R12.13 (acts on `C_interrupt`). The count of interruptions per hour imposed on the three nearest workstations by the room's traffic and noise shall be measured before and after occupation, and the delta shall be recorded as a cost of the room.

R12.14 (acts on `N`, `u`). Illuminance shall vary across the room by at least a factor of three between its brightest and dimmest usable working positions, so that lighting constitutes a real choice. Chapter 16 gives absolute ranges.

R12.15 (acts on `S`). Where a room exceeds 12 intended occupants, it shall be sub-divided into sub-volumes of not more than 6, by any means that establishes a distinct acoustic and visual sub-field. Chapter 13 gives non-wall methods; Chapter 14 gives the acoustic criterion.

R12.16 (acts on `eta`). A consumable point (water at minimum, hot drink preferred) shall be within 15 m and shall have a dwell of not less than 30 seconds per use.

R12.17 (acts on `C_exclusion`). The acoustic requirements of Chapter 14 are a precondition of this chapter's requirements, not an addition to them. A room failing Chapter 14 §14.5 shall not be certified against this chapter regardless of its other properties.

R12.18 (acts on `Q`). The room shall carry a visible, physical log of what it has produced — decisions taken, agreements standing. This is the room's `Q` made legible and it is the only defence against the Leaky Vessel failure mode.

R12.19 (ethics, CANON §5). No requirement in this chapter shall be implemented in a way that removes an occupant's ability to decline participation, to leave, or to be unobserved. Seating that cannot be exited, geometry that cannot be avoided, and observation that cannot be opted out of are out of scope for this manual.

R12.20 (measurement integrity). Any published figure for this room's `eta`, `delta`, `u` or `N` shall carry the date, the method, and the observer. Chapter 24's standard applies: **justify that the order proxy is not merely a repackaging of the input.**

The Measurement Worksheet

Reproduce this on one side of A4 per room per survey. Everything on it is measurable in one week by one person with a tape measure, a phone, and a clipboard.

```
=============================================================================
 ROOM SURVEY WORKSHEET                          Standard Manual, Ch. 12
=============================================================================
 Room ID ................  Building ..........  Floor ................
 Surveyor ...............  Survey week from ......... to ...........
 Intended function ..................................................

-----------------------------------------------------------------------------
 A. PHYSICAL INVENTORY                                        (one pass)
-----------------------------------------------------------------------------
 A1  Floor area, usable                            ............ m^2
 A2  Volume (area x mean clear height)             ............ m^3
 A3  Clear height                                  ............ m
 A4  Total surface area (floor+ceiling+walls)      ............ m^2
 A5  Glazed area                                   ............ m^2
 A6  Number of entrances                           ............
 A7  Longest internal sightline                    ............ m
 A8  Linear metres of sittable surface             ............ m
 A9  A8 / A1  (Whyte ratio; require >= 0.109)      ............ m/m^2
 A10 Movable seats / total seats                   ........./........
 A11 Seat height range measured                    ...... to ...... mm
 A12 Reverberation time RT60 at 500 Hz, occupied   ............ s   [Ch14]
 A13 Background noise level, unoccupied           ............ dBA [Ch14]
 A14 Illuminance max / min at working positions    ......./...... lux
 A15 Median door-to-door access time               ............ min

-----------------------------------------------------------------------------
 B. CATEGORY AUDIT  -> N                          (§12.10, strike rule)
-----------------------------------------------------------------------------
  #  posture         purpose           condition        available?
  1  ..............  ..............    ..............   Y / N
  2  ..............  ..............    ..............   Y / N
  3  ..............  ..............    ..............   Y / N
  4  ..............  ..............    ..............   Y / N
  5  ..............  ..............    ..............   Y / N
```

```
 6  ..............  .................   ..............   Y / N
    (strike duplicates and unavailable rows)
B1  N (surviving distinct categories)           .............
B2  ln N  (per-interaction ceiling)             ............. nats

-------------------------------------------------------------------------
C. ACTIVITY COUNT   (Gehl)        5 fixed times/day x 5 days = 25 obs
-------------------------------------------------------------------------
       NEC = necessary   OPT = optional   SOC = social
  day    NEC    OPT    SOC     OAR = OPT/(NEC+OPT)
  Mon   .....  .....  .....    ............
  Tue   .....  .....  .....    ............
  Wed   .....  .....  .....    ............
  Thu   .....  .....  .....    ............
  Fri   .....  .....  .....    ............
  TOT   .....  .....  .....    ............
C1  OAR (week)                                  .............
C2  SOC / OPT                                   .............
C3  Peak simultaneous occupancy                 .............
C4  Median dwell time, seated                   ............. min

-------------------------------------------------------------------------
D. CONTACT-PAIR AUDIT  -> u proxy                          (§12.10)
-------------------------------------------------------------------------
  pair  minutes  same team?  same discipline?  same lineage?  same line?
  ....  .......  Y/N         Y/N               Y/N            Y/N
  ....  .......  Y/N         Y/N               Y/N            Y/N
  ....  .......  Y/N         Y/N               Y/N            Y/N
  ....  .......  Y/N         Y/N               Y/N            Y/N
  (continue on reverse)
D1  Total conversational minutes observed       ............. min
D2  Minutes between pairs differing on >= 2 axes ............ min
D3  p_cross = D2 / D1   (badged u proxy)        .............
D4  N_eff = 1 + (B1 - 1) x D3                   .............
D5  F = ln(D4)                                  ............. nats
D6  F / B2  (fraction of ceiling reached)       .............

-------------------------------------------------------------------------
E. STORAGE TERMS
-------------------------------------------------------------------------
E1  Q proxy chosen (name it) ........................................
E2  Q_0 at survey start                          .............
E3  Q_1 after one zero-input period              .............
E4  delta = 1 - E3/E2                            ............. /period
E5  half-life = ln2 / (-ln(1-E4))                ............. periods
E6  S proxy chosen (name it) ........................................
E7  S per period                                 .............
E8  attributable dQ per period                   .............
```

```
 E9  eta_hat = E8 / E7                                  .............

-------------------------------------------------------------------------
 F. COST TALLY  -> C                                              (§12.8)
-------------------------------------------------------------------------
 F1  C_access     (min/session x sessions/period) .............
 F2  C_setup      (min/session x sessions/period) .............
 F3  C_effort     (Ch14 vocal effort dB; hoarseness reports) ..........
 F4  C_interrupt  (interruptions/hr at 3 nearest workstations) ........
 F5  C_exclusion  (Ch14 §14.7 ln N loss)          ............. nats
 F6  C_maint      (hours/period)                  .............
 F7  C_total, in the units of E7/E9 (state conversion) ................

-------------------------------------------------------------------------
 G. HEADROOM
-------------------------------------------------------------------------
 G1  eta ln N   = E9 x B2                         .............
 G2  delta Q    = E4 x E2                         .............
 G3  C          = F7                              .............
 G4  h = G1 - G2 - G3                             .............
      h < 0  UNSUSTAINABLE -> cut C, raise eta, or shed Q. Do not build.
      h = 0  MAINTENANCE ONLY -> hold.
      h > 0  GROWTH AVAILABLE -> r_max = h.

-------------------------------------------------------------------------
 H. COMPLIANCE
-------------------------------------------------------------------------
   R12.1 ... [ ]   R12.6 ... [ ]   R12.11 .. [ ]   R12.16 .. [ ]
   R12.2 ... [ ]   R12.7 ... [ ]   R12.12 .. [ ]   R12.17 .. [ ]
   R12.3 ... [ ]   R12.8 ... [ ]   R12.13 .. [ ]   R12.18 .. [ ]
   R12.4 ... [ ]   R12.9 ... [ ]   R12.14 .. [ ]   R12.19 .. [ ]
   R12.5 ... [ ]   R12.10 .. [ ]   R12.15 .. [ ]   R12.20 .. [ ]

 Surveyor signature .................  Date .......................
 Proxy justification (Ch24: Y is not a repackaging of U):
 ..........................................................................
 ..........................................................................
==========================================================================
```

The Failure Modes

Each entry gives the name, the clinical description, the damaged parameter, the diagnostic signature, the detecting measurement, and the fix. These entries feed Chapter 27's catalogue.

F12.1 — The Corridor With Chairs

Damages: `eta`, `S`. **Description.** A room whose occupancy is entirely necessary activity. People pass through, or attend because they were required to, and nobody is ever there by choice. **Signature.** OAR below 0.12. Median seated dwell under 4 minutes. Peak occupancy coincides exactly with scheduled events and drops to zero between them. **Detects with.** Worksheet section C. **Fix.** Triangulate (R12.7) and seat (R12.1, R12.2). A consumable point with real dwell is the single most reliable intervention. Expect OAR to roughly double within four weeks or the diagnosis was wrong.

F12.2 — The Beautiful Vessel

Damages: `eta`. **Description.** High optional activity, near-zero conversion. People linger happily; nothing becomes durable. The Leaky Vessel of Chapter 4, in its architectural form. This is the failure that falsifies Postulate 12.2 locally, and it is common in exactly the spaces that win design awards. **Signature.** OAR above 0.30 with attributable `dQ` per period at or near zero. Occupants describe the room warmly and cannot name anything that came out of it. **Detects with.** Worksheet E8/E9 measured *independently of* C1. Never infer `eta` from OAR alone. **Fix.** R12.8 (archived writable surface) and R12.18 (visible production log). The room needs an outlet, not more amenity.

F12.3 — The Single-Category Room

Damages: `N`, and therefore the ceiling `ln N`. **Description.** One posture, one purpose, one condition. The standard meeting room. `N = 1`, `ln N = 0` in the use-category dimension. **Signature.** Category audit strikes every row but one. Occupants who need a different mode leave the building rather than move within it. **Detects with.** Worksheet section B. **Fix.** Add categories, not area. A window ledge at 430 mm, a standing rail, a dimmable corner and one movable low table can take a room from `N = 1` to `N = 4` for the price of the ledge. Note the concavity: **going from one category to two is the largest single gain the room will ever record.**

F12.4 — The Public-Range Meeting

Damages: `u`, `C`. **Description.** Geometry places most pairs beyond 3.6 m. The channel degrades from conversation to announcement, and announcement carries no disagreement. **Signature.** Longest internal sightline in use exceeds 4 m for a group intended as one conversation. Speaking-turn distribution collapses onto two or three people (Chapter 14 §14.7 quantifies this). Participants at the far end report "I couldn't get in". **Detects with.** Worksheet A7 and D; the turn-share measurement of Chapter 14. **Fix.** Reduce the table, not the group; or split the group. R12.5 and R12.15.

F12.5 — The Nightly Reset

Damages: `delta` (drives it to 1 for the arrangement class of `Q`). **Description.** Every persistent physical state the occupants establish is erased between sessions by cleaning, by the next booking, or by a tidiness policy. The room can never accumulate convention. **Signature.** Occupants spend the first minutes of every session re-establishing the same arrangement. Whiteboards are

found blank. Nothing is ever where it was left. **Detects with.** Worksheet E2/E3/E4 with the arrangement as the `Q` proxy. **Fix.** R12.9 floor markings and R12.10 a written reset procedure owned by the last occupant, not by cleaning. Convert the convention from a memory into a marking.

F12.6 — The Welded Room

Damages: learning; the mirror image of F12.5. **Description.** `delta` near zero on everything. Nothing can be moved, so nothing can be improved. The room's arrangement is optimal for a use that ended in 2014. **Signature.** Movable seat fraction near zero (worksheet A10). Occupants have visible workarounds: sitting sideways, standing to see, using the floor. **Detects with.** A10, plus the trace instrument — look at the workarounds, they are the room's own bug report. **Fix.** Unweld. R12.2. Accept some `delta` in exchange for the ability to search.

F12.7 — The Externalised Room

Damages: `C_interrupt`, borne by non-users. **Description.** The room's costs fall on people outside it, who have no vote on its use. Classic commons failure in the Ostrom sense (Chapter 21): the boundary of the resource excludes some of those who bear its cost. **Signature.** Interruption rate at adjacent workstations rises measurably with room bookings. Adjacent occupants have adopted headphones, moved, or stopped using their desks. **Detects with.** Worksheet F4, measured before and after. **Fix.** Acoustic separation (Chapter 14), traffic re-routing (Chapter 15), and governance: extend the decision boundary to include the affected. **The engineering fix and the governance fix are both required; either alone regresses.**

F12.8 — The Uncounted Exclusion

Damages: `N`, `u`, and the credibility of every measurement you take. **Description.** The room suppresses certain contributions — acoustically, proxemically, or by sightline — and because a suppressed contribution leaves no trace, the room's own metrics look fine. The survey reports a healthy OAR from the people the room did not exclude. **Signature.** Turn-share concentration. Systematic absence of a subgroup. Consistent late-cycle objections from people who "should have raised it earlier" and could not. **Detects with.** Chapter 14 §14.7 computes the exclusion from physics rather than from report. This is the only reliable route, because **self-report cannot detect a contribution that was never attempted. Fix.** Chapter 14's specification in full, plus Chapter 17's protocol layer (silent-write-first, round-robin) as a backstop. R12.17 makes the acoustic requirement a precondition.

F12.9 — The Amenity Substitution

Damages: honesty of the headroom calculation. **Description.** Visible amenity is installed and reported as if it were measured improvement. Plants, art, a coffee machine, and a rebrand, with no before/after measurement and no change to `eta`, `delta`, or `C`. **Signature.** The intervention has a photograph and no numbers. The evaluation is a satisfaction survey administered once, after. **Detects with.** Chapter 25's standard: preregistered before-measurement, held-out after-measurement, and the shuffled-input control. **Fix.** Measure before you buy. If the worksheet was not filled in before the intervention, the intervention has not been evaluated, whatever the survey says.

F12.10 — The Observed Room

Damages: trust; violates CANON §5.4 and R12.19. **Description.** Measurement conducted in a way that removes occupants' ability to be unobserved. Continuous individual-level tracking, badge telemetry tied to identity, cameras without notice. **Signature.** Occupants change behaviour on learning of the instrumentation, which also destroys the measurement. **Detects with.** Ask. This is the one failure mode where the intercept instrument is the right one. **Fix.** Aggregate counts, posted notice, defined retention, opt-out that costs nothing. **A measurement that cannot be declined is not permitted by this manual, and it is also bad instrumentation, because it measures the observation and not the room.**

Worked Example

The subject

A departmental common room, 120 m^2, clear height 2.8 m, volume 336 m^3, on the second floor of a research building. It serves roughly 90 people across four teams. It contains a kitchenette with a kettle and a fast drip machine, six four-person tables with 24 stacking chairs, a notice board, and a vending machine. Floor is sealed screed; ceiling is exposed soffit with services; two walls are glazed. It was refurbished eighteen months ago and the department describes it as "underused despite being lovely."

Step 1 — Physical inventory (worksheet A)

Item	Measured
A1 Floor area	120 m^2
A2 Volume	336 m^3
A8 Linear sittable metres	6.0 m (six tables x 4 chairs, chair width ~0.45 m, but only the chairs are sittable; the window ledges are 220 mm deep and therefore not sittable under R12.3)
A9 Whyte ratio	6.0 / 120 = 0.050 m/m^2
A10 Movable / total seats	24 / 24, but all are identical stacking chairs at identical tables
A15 Median access time	1.5 min

Requirement R12.1 demands 120 x 0.109 = 13.1 m of sittable edge. The room has 6.0 m. Deficit: 7.1 m. The room is at 46% of the Whyte floor.

Note the trap in A10. Every chair is movable, so the room scores 100% on movability, and yet it offers no seating *choice*, because all 24 seats are the same seat at the same table at the same height under the same light. **Movability without variety satisfies the letter of R12.2 and fails its purpose. R12.4's category audit is the check that catches it.**

Step 2 — Category audit (worksheet B)

#	Posture	Purpose	Condition	Available?
1	Sitting upright, standard chair	Eating / pair talk	Bright, live, open	Y
2	Sitting upright, standard chair	Solo focus	Bright, live, open	Y — but identical to 1 in every respect that matters to a user
3	Standing	Queue at kettle	Bright, live, open	Y
4	Sitting reclined	Anything	–	N (no such seat exists)
5	Sitting, enclosed / backed	Small-group deliberation	Quiet	N
6	Perching	Brief transit stop	–	N (ledges 220 mm)

Rows 4, 5 and 6 struck as unavailable. Row 2 struck as a duplicate of row 1 under the §12.10 strike rule. Surviving categories: **2**.

```
N   = 2
ln N = 0.6931 nats
```

The room's per-interaction ceiling in the use-category dimension is 0.693 nats — 33% of what a six-category room would offer (`ln 6 = 1.792`).

Step 3 — Activity count (worksheet C)

Five fixed observation times per day (09:30, 11:00, 12:45, 14:30, 16:00), five days, marked on a plan. Counts are person-observations summed over the week.

Day	NEC	OPT	SOC	OAR
Mon	41	7	3	0.146
Tue	38	5	2	0.116
Wed	44	9	4	0.170
Thu	36	6	2	0.143
Fri	52	21	11	0.288
Total	**211**	**48**	**22**	**0.185**

`SOC / OPT = 22 / 48 = 0.458`.

Read the Friday anomaly. Friday is the day the department buys pastries. **The single largest determinant of optional activity in this room is whether there is food in it, and it is not close.** That is a Whyte finding reproduced, by counting, in one week, by one person.

Necessary activity is 211 observations — people come to get coffee and leave. OAR of 0.185 is a corridor-with-chairs signature (F12.1) with a food-driven Friday exception.

Median seated dwell, timed from the time-lapse: **4.1 minutes**.

Step 4 — Contact-pair audit (worksheet D)

Conversational minutes observed over the week: 412. Minutes between pairs differing on at least two of {team, discipline, training lineage, reporting line}: 96.

```
p_cross = 96 / 412 = 0.233     (badged: a crude u proxy, see §12.10)
N_eff   = 1 + (2 − 1) x 0.233 = 1.233
F       = ln(1.233) = 0.2095 nats
F / ln N = 0.2095 / 0.6931 = 0.302
```

The room reaches 30% of its own (already low) ceiling. Note that N = 2 caps this hard: even at u = 1 the room could deliver only 0.693 nats. **Raising u in this room cannot help much, because N is the binding constraint.** This is the diagnosis the arithmetic is for, and it is the opposite of the intervention the department had planned (a "mixer" social event).

Step 5 — Storage terms (worksheet E)

Q proxy, declared: **standing cross-team working agreements attributable to contact in this room** — a countable list, maintained by the department administrator, with a named origin for each entry.

Decay measurement: the room was closed for three weeks over the summer shutdown, giving a genuine zero-input period. Agreements standing before: 9. Agreements still active after (three periods, period = one week… the shutdown gives us a three-period decay):

```
Q_0 = 9,  Q_3 = 7
(1 − delta)^3 = 7/9 = 0.7778
1 − delta = 0.7778^(1/3) = 0.9192
delta = 0.0808 per week
half-life = ln2 / (−ln(0.9192)) = 0.6931 / 0.08418 = 8.23 weeks
```

S proxy, declared: **observed social contacts per week** (SOC count) = 22.

Attributable dQ per week, from the administrator's log: **1.0** new standing agreement per week originating in this room.

```
eta_hat = 1.0 / 22 = 0.0455
```

Step 6 — Cost tally (worksheet F)

Converting a cost tally into the same units as eta S requires a declared conversion, and the declaration is the honest part. Here we express everything as *equivalent lost agreements per week*, the same unit as Q, and state the conversion openly.

Component	Raw measure	Converted (agreements/week equivalent)	Basis of conversion
C_access	1.5 min x 60 visits/wk = 90 min	0.10	15 min of aggregate time judged equal to 1/60 of an agreement's setup effort
C_setup	negligible (no equipment)	0.00	–
C_effort	RT60 measured at 1.6 s (Chapter 14), background 47 dBA	0.55	Chapter 14 §14.7 exclusion arithmetic
C_interrupt	0 (room is acoustically separate)	0.00	–
C_exclusion	included in C_effort above	–	do not double-count
C_maint	2 h/wk cleaning	0.05	–
C_total		**0.70**	

The conversion is a modelling choice and it is the weakest link in this calculation. State it, sensitivity-test it, and do not present the resulting h to three decimal places as if it were measured. The correct use of this number is comparative — before versus after — not absolute.

Step 7 — Headroom (worksheet G)

```
eta ln N  = 0.0455 x 0.6931 = 0.0315
delta Q   = 0.0808 x 9      = 0.7272
C         =                   0.70
h = 0.0315 − 0.7272 − 0.70  = −1.3957
```

h < 0. Unsustainable in this model. No arrangement of people in this room prevents the stock of cross-team agreements from decaying. The department's plan — a monthly mixer — would add S, and S enters the budget multiplied by eta = 0.0455. To close a gap of 1.396 by adding S alone would require an implausible increase, and in any case S is capped: eta ln N = 0.0315 is the *maximum* the term can ever reach at N = 2, no matter how much interaction occurs.

This is Corollary 3.4 doing its job. It tells you, before you spend anything, that the planned intervention cannot work.

Step 8 — The intervention, chosen by the arithmetic

The binding constraints, in order of magnitude: delta Q = 0.727, C = 0.70, and an eta ln N term so small it is almost irrelevant. Three moves follow directly.

Move 1 — Raise N from 2 to 5. This is the cheapest large move available because N = 2 is capping everything.

- Deepen the two window ledges from 220 mm to 450 mm over 9 linear metres (perching, transit stop, bright, solo). New category.
- Install one high table with four stools and a rail (standing, pair work, live). New category.
- Install two backed banquettes forming a pair of 4-person alcoves against the solid wall, with local dimmable light (sitting enclosed, small-group deliberation, quiet). New category.
- Retain the existing tables (category 1) and the kettle queue (category 3).

```
N: 2 -> 5          ln N: 0.6931 -> 1.6094
```

Move 2 — Cut `C` by 0.55. The entire `C_effort` term is acoustic. Chapter 14's treatment for a 336 m^3 room brings RT60 from 1.6 s to under 0.7 s and background from 47 dBA to 40 dBA. Cost is one absorptive ceiling raft array and 18 m^2 of wall panel.

```
C: 0.70 -> 0.15
```

Move 3 — Shed `Q`. Of the nine standing agreements, four are dormant: they name people who have left, or processes that ended. They are consuming `delta Q` — 0.323 agreements/week of maintenance budget — and returning nothing. Retire them formally.

```
Q: 9 -> 5          delta Q: 0.727 -> 0.404
```

Move 3 is the one nobody proposes, and Chapter 6 named it: **shedding stored order to regain headroom is a legitimate and underused move.**

Alongside these, R12.1 is satisfied: the deepened ledges (9 m) plus high table (4 stools ~1.8 m) plus banquettes (8 seats ~3.6 m) plus existing chairs (6.0 m) gives 20.4 m of sittable edge against the 13.1 m requirement. R12.7 is satisfied by relocating the notice board to the kettle queue and converting it to a live, changing display. R12.16 is satisfied by replacing the fast drip machine with a bean-to-cup unit whose cycle is 55 seconds — deliberately slower, per §12.6.

Step 9 — Predicted outcome and what would falsify it

Recomputing headroom with `eta` held at its measured value (we predict `eta` will rise, but we refuse to assume it):

```
eta ln N  = 0.0455 x 1.6094 = 0.0732
delta Q   = 0.0808 x 5      = 0.4040
C         =                   0.15
h = 0.0732 - 0.4040 - 0.15  = -0.4808
```

Still negative. **Say so.** The three moves close the gap from -1.396 to -0.481 — 66% of the way — and the room is still, in this model, not able to grow its stock of agreements.

The remaining gap is dominated by `delta Q = 0.404` against an `eta ln N` of only 0.073. The structural reading: **this room's `eta` is an order of magnitude too small for the `Q` it is being asked to hold.** At `eta = 0.0455`, sustaining even `Q = 5` at `delta = 0.081` would require `eta ln N >= 0.404`, hence `ln N >= 8.88`, hence `N >= 7200` categories. Absurd. The room cannot be the mechanism that sustains these agreements.

That is the finding, and it is more valuable than a success would have been. **The agreements are not sustained by the room; they must be sustained by an artefact.** Chapter 4's answer applies: the agreements need a written register with a named owner and a review cadence, which raises `eta` at the *organisational* level and lowers `delta` by moving `Q` out of memory and into a

document. The room's job is to generate contact — which the three moves genuinely improve — not to store the result.

Recompute once more, with the register in place. The register is measured, after twelve weeks, to have moved `delta` from 0.0808 to 0.030 per week (agreements now decay only when genuinely obsolete, not when forgotten) and `eta_hat` from 0.0455 to 0.098 (contacts now have somewhere to land):

```
eta ln N  = 0.098 x 1.6094  = 0.1577
delta Q   = 0.030 x 5       = 0.1500
C         =                   0.15
h = 0.1577 − 0.1500 − 0.15  = −0.1423
```

Closer. Still negative, and honestly so. Adding the observed post-intervention `SOC` rise (22 to 56 contacts/week — see below) does not change h, because h is a budget on the *ceiling*, not on the realised input; but it does change whether the realised `S_t` approaches that ceiling.

The forecast the department should be given is therefore: **this room, treated, will roughly double its optional activity and roughly double its `eta`; it will not on its own make the department's stock of cross-team agreements grow, and anyone who promises that is not reading the arithmetic.**

Preregistered predictions, per Chapter 25:

#	Prediction	Measured how	Falsified if
P1	OAR rises from 0.185 to >= 0.30 within 8 weeks	Worksheet C, same 5 times, same 5 days	OAR < 0.25 at week 8
P2	Median seated dwell rises from 4.1 to >= 8 min	Time-lapse	< 6 min at week 8
P3	`SOC` count at least doubles	Worksheet C	< 1.5x at week 8
P4	`p_cross` rises (the alcoves and rail mix teams)	Worksheet D	p_cross < 0.25 at week 8
P5	RT60 at 500 Hz occupied <= 0.70 s	Chapter 14 method	> 0.80 s at commissioning
P6	h improves by at least 0.8	Worksheet G	improvement < 0.5
P7	The Friday food effect persists but shrinks as a fraction of weekly OPT	Worksheet C	Friday remains > 40% of weekly OPT

P7 is the control that matters. **If all the gain is still on pastry day, the room was not fixed; the catering was.** That is the shuffled-input control of Chapter 25 in its most concrete possible form.

Step 10 — Measured outcome at week 8

Day	NEC	OPT	SOC	OAR
Mon	43	19	9	0.306
Tue	40	17	8	0.298
Wed	45	24	13	0.348
Thu	38	18	9	0.321
Fri	50	31	17	0.383
Total	**216**	**109**	**56**	**0.335**

```
OAR:  0.185 -> 0.335   (x1.81)     P1 PASSES (>= 0.30)
SOC:  22    -> 56      (x2.55)     P3 PASSES
Friday share of weekly OPT: 21/48 = 43.8%  ->  31/109 = 28.4%
                                               P7 PASSES (< 40%)
eta_hat: 1.0/22 = 0.0455  ->  4.0/56 = 0.0714   (x1.57)
```

Necessary activity is essentially unchanged (211 to 216), which is exactly what Gehl's taxonomy predicts: **necessary activity does not respond to quality, and its stability across the intervention is evidence that the observation protocol was not itself the cause.** That stability is the strongest internal validity check the worksheet provides, and it is free.

The room now reaches `N = 5`, `p_cross` measured at 0.31, hence `N_eff = 1 + 4 x 0.31 = 2.24`, `F = 0.807` nats against a ceiling of 1.609 — 50% of ceiling, up from 30%.

What this example demonstrates

1. **The worksheet found the binding constraint in one week, and it was not the one the department believed.** They believed the problem was social; the arithmetic said `N = 2` and `delta Q` were binding.
2. **The intervention chosen by the arithmetic was cheaper than the one that had been planned** — deeper ledges, a rail, two banquettes and a slower coffee machine, against a recurring events programme.
3. **The honest answer was partly negative and was reported as such.** h improved substantially and remained below zero, and the correct conclusion was that the room is a contact generator, not a storage device.
4. **The Friday control did the real work.** Without P7, a department could have claimed the refurbishment succeeded when all it had done was buy pastries.

Chapter 13 takes the sub-division problem this example raised — how to get five categories and four alcoves into 120 m^2 without building walls — and solves it with geometry. Chapter 14 takes the `C_effort` term, which was the largest single cost line in this survey, and reduces it to a computation.

Chapter 13 . Fractal and Parabolic Space

Chapter 12 ended with a problem it could not solve. A 120 square metre room needed five distinct use-categories, four small-group sub-volumes, and twenty metres of sittable edge, and it needed them without walls — because walls consume the floor area that made the room worth having, and because a wall is the highest-`delta`-resistance, highest-`C` intervention in the catalogue.

This chapter is the geometry of solving that. Its subject is **shape that performs a function without consuming power, and shape that performs the same function at several scales at once.**

Those are two different tricks and they have two different mathematics. The first is the parabola, whose reflective property is a theorem of elementary geometry and which was built, in concrete, at sixty-one metres, to detect aircraft before radar existed. The second is self-similarity, whose arithmetic is the fractal dimension and whose commercial vindication is sitting in every mobile handset on earth. Both are real. Both are routinely over-claimed. This chapter states exactly what each one buys and, more carefully, what it does not.

The Law

What is proven

Theorem 13.1 (the reflective property of the parabola) — PROVEN, elementary geometry. For the parabola `y = x^2/(4f)` with focus `F = (0, f)` and directrix `y = −f`, every ray parallel to the axis of symmetry reflects, at the surface, through `F`.

Theorem 13.2 (the focus-directrix property) — PROVEN. For any point `P = (a, a^2/(4f))` on that parabola, `|PF| = y_P + f`, which is exactly the perpendicular distance from `P` to the directrix.

Corollary 13.3 (path-length equality — the load-bearing one) — PROVEN. Every axis-parallel ray descending from a common height `h` and reflecting to the focus travels the same total distance `h + f`, independent of where on the surface it struck.

Corollary 13.3 is why a parabola has *gain* and a merely-curved surface does not. Convergence without path-length equality produces arrival, but arrival out of phase, which is to say partial cancellation. Convergence *with* path-length equality produces coherent addition. **The parabola is not a funnel. It is a phase-correcting device, and the phase correction is done by the shape itself, at no cost, forever.**

Theorem 13.4 (similarity dimension) — PROVEN, by definition of self-similarity. A set composed of `m` copies of itself, each scaled by `1/r`, has similarity dimension `D = ln m / ln r`.

Theorem 13.5 (crenellation multiplier) — PROVEN, trivially, by summing segments. A boundary replaced by a square-wave of tooth depth `d` and half-period `w` has its length multiplied by exactly `1 + d/w`.

What is postulated

POSTULATE 13.1 (badged): hyperrecursive symmetry. The storage law `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` applies at every scale of enclosure — desk, alcove, room,

wing, campus — with the same *form* and different *parameters*. CANON §1.4 states the honesty condition on this and it is repeated here without softening: **this is a claim about model form, not about physics.** There is no theorem that a campus obeys the same recursion as a desk. There is a design decision to model both that way, and it earns its keep only when the parameters at each scale are separately measured. §13.10 measures them.

POSTULATE 13.2 (badged): a self-similar boundary buys multi-scale social response the way a self-similar conductor buys multi-band electrical response. The analogy is structural and it is suggestive. It is **not** a derivation. Antenna self-similarity produces resonances at predictable frequency ratios because Maxwell's equations are scale-invariant in the relevant sense. Nothing comparable is proven for social occupancy. §13.8 states exactly how far the analogy carries and where it stops.

The honesty clause this chapter cannot do without

Fractal geometry performance is strongly geometry-dependent and requires modelling. It cannot be assumed from the concept.

This deserves its own paragraph because it is the single most abused idea in contemporary design writing. Two structures with identical fractal dimension can have wildly different acoustic, electromagnetic, and occupancy behaviour. Dimension is one number summarising a scaling relationship; it does not determine the response. In the antenna literature — which is the most rigorous body of work on engineered fractals — the multiband behaviour of self-similar radiators is well established, while **the stronger claim that fractality *per se* is optimal is contested, with several studies finding conventional designs of the same overall size performing comparably.**

The rule for this manual: **you may use a fractal argument to generate a candidate geometry. You may not use it to certify one. Certification requires a model or a measurement of that specific shape.** Every requirement in §13.13 that invokes self-similarity carries a modelling obligation attached to it.

The Engineering

13.1 Geometry that does work without power

Return to the cost term. `C_t` in the storage law is dissipation: what the system spends to keep running. A powered device that performs a spatial function — an amplifier, a fan, a display, an active noise system — enters the budget as a permanent `C` line. It needs electricity, maintenance, replacement, a technician, and a failure mode.

A shape enters the budget once, at construction, and then contributes to `C` only through cleaning and eventual repair.

```
Powered function:  C_t = c_energy + c_maint + c_replace + c_failure  , every period
Geometric function: C_t = c_clean                                    , every period
                      plus a one-time capital charge at t = 0
```

This is the whole argument for Part III. A shape is a function with a `C` of nearly zero and a `delta` set by material durability rather than by anyone's attention. It is Weiser's disappearing

technology (CANON §3) in its most literal form: a correctly shaped alcove performs its work without being noticed, without being switched on, and without anyone remembering that it is there.

Star and Ruhleder's observation applies with force: infrastructure becomes visible only on breakdown. A shape's breakdown mode is slow (wear, obstruction, misuse) rather than sudden (power cut, firmware, obsolescence), which makes it a low-variance element in a system's cost profile. Chapter 3's treatment of `delta` gives the general form; here the specific point is that **stone, concrete, hardwood and plaster have `delta` values measured in decades, and almost nothing else in an organisation does.**

13.2 The parabola, derived

Take the parabola in its focal form:

```
y = x^2 / (4f)            focus  F = (0, f)
                          directrix  y = −f
                          dy/dx = x / (2f)
```

Step 1 — the focus-directrix property (Theorem 13.2). Let `P = (a, a^2/(4f))`. Then

```
|PF|^2 = a^2 + ( a^2/(4f) − f )^2
       = a^2 + a^4/(16 f^2) − 2f ∗ a^2/(4f) + f^2
       = a^2 + a^4/(16 f^2) − a^2/2 + f^2
       = a^2/2 + a^4/(16 f^2) + f^2

( a^2/(4f) + f )^2 = a^4/(16 f^2) + 2f ∗ a^2/(4f) + f^2
                   = a^4/(16 f^2) + a^2/2 + f^2
```

The two expressions are identical, so `|PF| = a^2/(4f) + f = y_P + f`, which is the vertical distance from P down to the directrix `y = −f`. **The distance to the focus equals the distance to the directrix, at every point on the curve. That is the parabola's definition, recovered from its equation.**

Step 2 — the reflective property (Theorem 13.1). A ray descending parallel to the axis has direction `d = (0, −1)`. The surface normal at P is obtained from the gradient of `y − x^2/(4f)`, giving `n proportional to (−a/(2f), 1)`, normalised. The reflected direction is

```
r = d − 2 (d . n_hat) n_hat
```

Carrying this out numerically for `f = 1.0 m` and checking where the reflected ray crosses the focal height `y = f`:

a (m)	y_P (m)	slope a/2f	tangent angle (deg)	x-crossing at y = f
0.25	0.0156	0.1250	7.13	0.000000
0.50	0.0625	0.2500	14.04	0.000000
0.75	0.1406	0.3750	20.56	0.000000
1.00	0.2500	0.5000	26.57	0.000000
1.50	0.5625	0.7500	36.87	0.000000

a (m)	y_P (m)	slope a/2f	tangent angle (deg)	x-crossing at y = f
2.00	1.0000	1.0000	45.00	(reflected ray is exactly horizontal here; y_P = f)
3.00	2.2500	1.5000	56.31	0.000000

Every reflected ray crosses y = f at x = 0. The a = 2f row is the degenerate case: at that point the reflected ray leaves horizontally, travelling straight along y = f to the focus. That is the rim-angle-90-degree point and it is why f/D = 0.25 is the geometry where the rim subtends a full half-space at the focus.

Step 3 — path-length equality (Corollary 13.3). Let each ray begin at height h above the vertex, descend to P, then travel to F:

```
total = (h − y_P) + |PF| = (h − y_P) + (y_P + f) = h + f
```

The y_P cancels. Numerically, with f = 1.0 and h = 5.0:

a (m)	y_P	|PF|	(h − y_P) + |PF|
0.0	0.0000	1.0000	6.0000
0.5	0.0625	1.0625	6.0000
1.0	0.2500	1.2500	6.0000
1.5	0.5625	1.5625	6.0000
2.0	1.0000	2.0000	6.0000
3.0	2.2500	3.2500	6.0000

Identical to four decimal places at every aperture radius, because it is identically h + f. This is the theorem doing the work that an electronic phase-alignment network would otherwise have to do, and it does it with zero parts and zero power.

The shape IS the function. There is no computation, no calibration, no control loop. The function is stored in the boundary condition.

13.3 Aperture, depth, and f-ratio

For a dish of aperture diameter D and focal length f, the rim sits at x = D/2, so:

```
depth  =  (D/2)^2 / (4f)  =  D^2 / (16 f)
f/D    =  the "f-ratio"; smaller means deeper
```

D (m)	f (m)	f/D	depth (m)	rim angle at focus (deg)
0.50	0.15	0.300	0.104	79.6
1.00	0.30	0.300	0.208	79.6
2.00	0.50	0.250	0.500	90.0
6.10	1.50	0.246	1.550	90.9

D (m)	f (m)	f/D	depth (m)	rim angle at focus (deg)
9.10	2.30	0.253	2.250	89.4
61.00	15.00	0.246	15.504	90.9

The f/D around 0.25 recurs because at f/D = 0.25 the focus sits exactly in the plane of the rim — the shallowest geometry that still collects a full hemisphere at the focus, and therefore the most material-efficient. **A designer who converges on f/D near 0.25 has usually converged correctly and does not need to justify it further.**

Deeper dishes (f/D < 0.25) shield the focus from off-axis noise better, at the cost of material and of a focus buried inside the shell where a head does not comfortably fit. Shallower dishes (f/D > 0.3) put the focus out in the open where it is easy to use and easy to contaminate.

13.4 What a passive aperture actually buys, and its hard frequency floor

Here is where over-claiming begins, so here is the arithmetic.

Directivity. For a circular aperture of diameter D at wavelength lambda, the half-power beamwidth is approximately

```
theta ~ 70 * lambda / D    degrees      (standard aperture approximation)
```

Aperture gain. The theoretical on-axis gain, with aperture efficiency eta_ap:

```
G = (pi D / lambda)^2 * eta_ap
```

Take eta_ap = 0.5, a conventional and generous value for a real reflector. Speed of sound at 20 C is c = 331.3 + 0.606 T = 343.4 m/s; use 343 m/s throughout.

For a 1.2 m aperture — the largest that fits comfortably behind a seated head:

Frequency	lambda (m)	D/lambda	Beamwidth	Aperture gain (upper bound)
250 Hz	1.372	0.87	80.0 deg	5.8 dB
500 Hz	0.686	1.75	40.0 deg	11.8 dB
1000 Hz	0.343	3.50	20.0 deg	17.8 dB
2000 Hz	0.172	7.00	10.0 deg	23.8 dB
4000 Hz	0.086	13.99	5.0 deg	29.9 dB

The frequency floor. A reflector does essentially nothing when it is small compared with the wavelength; sound simply diffracts around it. The working rule is that useful directivity begins at about D = 2 lambda:

Aperture D	D = 1 lambda at	D = 2 lambda at (useful floor)	D = 4 lambda at
0.30 m	1143 Hz	2287 Hz	4573 Hz
0.50 m	686 Hz	1372 Hz	2744 Hz
0.60 m	572 Hz	1143 Hz	2287 Hz

Aperture `D`	`D = 1 lambda` at	`D = 2 lambda` at (useful floor)	`D = 4 lambda` at
0.90 m	381 Hz	762 Hz	1524 Hz
1.20 m	286 Hz	572 Hz	1143 Hz
2.00 m	172 Hz	343 Hz	686 Hz
3.00 m	114 Hz	229 Hz	457 Hz
6.10 m	56 Hz	113 Hz	225 Hz
9.10 m	38 Hz	75 Hz	225 Hz
61.00 m	5.6 Hz	11 Hz	22 Hz

Now place speech on that scale:

```
   adult fundamental          85 – 255 Hz
   vowel formants F1, F2      250 – 2500 Hz
  consonants, fricatives    2000 – 8000 Hz   <-- where intelligibility lives
```

A 0.6 m shell behind a chair has a floor of about 1143 Hz. It does nothing whatever to the voice's fundamental, little to the first formant, and real work only on the consonant band. That is not a failure; the consonant band is precisely where intelligibility is won or lost, and Chapter 14 will show that consonants are also exactly what a reverberant room destroys first. But it means the honest description of a small acoustic shell is **"a consonant-band directivity device", not "it makes you louder."** Anyone selling the second description is selling something else.

Three further honesty conditions on the gain figures above:

1. **The gain formula assumes a plane wave arriving on-axis from the far field.** A talker 1.5 m from a 1.2 m dish is nowhere near the far field of that aperture; the wavefront is curved across the aperture and the coherent addition is partial. Real near-field performance is substantially below the tabulated bound.
2. **`eta_ap = 0.5` is an assumption.** Surface accuracy, edge diffraction, blockage by the listener's own head at the focus, and absorption in the shell material all reduce it. Head blockage in particular is severe in listening shells, because the thing you are trying to feed is sitting in the aperture.
3. **These are single-frequency figures.** Speech is broadband; the perceived effect is an average weighted by the speech spectrum and by the ear's frequency response, and it will be less impressive than the 4 kHz row.

Rule: treat every computed aperture gain as an upper bound that will not be reached, and measure the built article.

13.5 Denge: the real precedent

Between roughly 1928 and the early 1930s, at Denge near Dungeness on the Kent coast, the British Air Ministry built a set of concrete acoustic mirrors as an early-warning system for incoming aircraft. The work was led by William Sansome Tucker. Three principal structures survive: a 20-foot (6.1 m) dish, a 30-foot (9.1 m) dish, and a 200-foot (61 m) curved wall. Similar mirrors were built elsewhere along the coast.

The operating principle is exactly Theorem 13.1. Aircraft engine noise arriving as a near-plane

wave from a distance is collected by the concrete surface and concentrated at the focus, where a microphone (in earlier installations, a listening trumpet and a human ear) picks it up. An operator determined bearing by moving the collector across the focal region. **There is no power in the acoustic path. The concrete does the work.**

What the geometry bought, in this chapter's terms:

`D`	Beamwidth at 250 Hz	Beamwidth at 500 Hz	Aperture gain at 250 Hz (bound)
6.10 m	15.7 deg	7.9 deg	19.9 dB
9.10 m	10.6 deg	5.3 deg	23.4 dB
61.00 m	1.6 deg	0.8 deg	46.9 dB

The 61 m wall is a strip mirror rather than a dish, so it focuses in one plane and the tabulated circular-aperture figures are indicative only — but the order of magnitude is right, and it explains why they built something that large: **they needed bearing resolution on a low-frequency source, and bearing resolution is set by `lambda/D` and by nothing else.** At 250 Hz you cannot get a two-degree beam out of anything smaller than a building.

Reported detection ranges were of the order of tens of kilometres in favourable conditions, yielding a warning of some minutes. Treat the specific range figures with caution — they varied enormously with wind, sea state, aircraft type and operator skill, and the surviving accounts are not consistent. **The engineering claim that is safe to make is the one the geometry guarantees: the mirrors provided real, passive, unpowered directional gain in the low-frequency band, and they worked.**

Why Denge was abandoned, and why the reason matters Radar — RDF, following Watson-Watt's demonstration in 1935 — rendered the mirrors obsolete within a few years. It is tempting to read this as the geometry failing. It is not what happened, and the actual reason is a lesson this manual needs.

The mirrors did not fail. Their input signal moved out from under them. Three things changed simultaneously:

1. **Aircraft got faster.** Warning time is range divided by closing speed. The mirrors' range was fixed by physics; as speeds rose, the minutes of warning shrank toward uselessness. The device was constant and the requirement moved.
2. **The acoustic environment got noisier.** Road traffic and increasing civil aviation raised the background against which the operator was listening. A passive device has no gain control and no way to reject a competing source that lies within its beam.
3. **Radar carried its own signal.** This is the decisive difference. A passive collector can only concentrate what happens to arrive. An active system chooses what to transmit, when, and at what power, and therefore controls its own signal-to-noise ratio.

The general principle, and it applies to every geometric intervention in this manual: passive geometry has no gain control and generates no signal of its own. It can only concentrate what arrives. A perfectly shaped alcove in an organisation with nothing to say will concentrate silence with excellent efficiency.

That is the correct scope statement for Part III. Geometry sets the coefficients. It does not supply `S`.

13.6 Where the parabola belongs in a building — and where it emphatically does not

It belongs at furniture and niche scale, aimed deliberately.

Application	Aperture	Frequency floor	What it actually does
Listening niche behind a seated head	1.0 - 1.4 m	~490 - 690 Hz	Raises consonant-band SNR for one seated person facing the room
Desk-scale shell	0.5 - 0.7 m	~980 - 1372 Hz	Consonant-band privacy shield; reduces spill both ways
Alcove back-shell (shallow, `f/D` ~0.35)	2.0 - 2.6 m	~264 - 343 Hz	Sub-divides an open floor acoustically without a wall
Outdoor listening seat / whisper pair	1.5 - 3.0 m	~229 - 457 Hz	Triangulating object (§12.6) that is also a demonstration

It does not belong as the shape of a room.

This is the most important negative in the chapter. **In a room, concave surfaces are usually a defect.** A domed ceiling, a curved end wall, a barrel vault, or a concave glazed façade concentrates sound at focal points, producing:

- **Hot spots**: positions where a talker is uncomfortably loud and positions two metres away where they are inaudible. This directly destroys the turn-taking equality that Chapter 14 is about.
- **Focused echoes**: a discrete, delayed, high-level reflection at a specific seat. Discrete echoes are far more damaging to intelligibility than an equal amount of diffuse reverberation, because the ear integrates diffuse energy and cannot separate a strong late arrival from the direct sound.
- **Position-dependent performance**, which is a governance problem as much as an acoustic one: a room where the best seat is acoustically privileged has installed a head of the table by accident. Chapter 17 explains why that is fatal.

Rule R13.9 below states it as a prohibition: no concave surface with a radius of curvature comparable to a room dimension shall face the occupied zone, unless it is diffusive, absorptive, or deliberately focused away from every seat. The parabola is a scalpel. It is not wallpaper.

13.7 Self-similarity, and what the dimension does and does not tell you

A set built from `m` copies of itself at scale `1/r` has similarity dimension `D = ln m / ln r`:

Set	m	1/r	D = ln m / ln r
Cantor set	2	3	0.6309
Koch curve	4	3	1.2619
Sierpinski gasket	3	2	1.5850
Sierpinski carpet	8	3	1.8928
Menger sponge	20	3	2.7268
Hilbert curve	4	2	2.0000

The Koch snowflake makes the useful property visible. Its perimeter grows without bound while its area converges:

iteration n	perimeter (unit-side triangle)	x original	area / A_0
0	3.0000	1.0000	1.00000
1	4.0000	1.3333	1.33333
2	5.3333	1.7778	1.48148
3	7.1111	2.3704	1.54733
4	9.4815	3.1605	1.57659
5	12.6420	4.2140	1.58960
6	16.8560	5.6187	1.59538
7	22.4746	7.4915	1.59794
8	29.9662	9.9887	1.59909

Area converges to exactly 8/5 = 1.6 times the original triangle. Perimeter multiplies by 4/3 every iteration and diverges.

This is the entire architectural point of fractals in one table: boundary is cheap, interior is expensive, and self-similarity is the technique for buying enormous amounts of the first with bounded amounts of the second.

Gehl's edge effect (Chapter 12 §12.2, Chapter 15 in full) says people occupy edges first. Whyte's sittable-space ratio says a space's use is governed by its sittable perimeter. Both are demands for boundary. A room's floor area is fixed by the lease; its boundary length is a design variable, and Theorem 13.5 prices it exactly.

What the dimension does not tell you. D = 1.2619 for the Koch curve is a statement about how its measured length scales with ruler size. It says nothing about:

- whether the shape scatters sound usefully at any particular frequency,
- whether people will sit in its recesses,
- whether the recesses are reachable, cleanable, or safe,
- whether the geometry resonates, focuses, or traps anything undesirable.

Two boundaries with identical D can behave completely differently. This is why Postulate 13.2 is badged and why every self-similarity requirement in §13.13 carries a modelling obligation.

13.8 Why fractal antennas are in handsets

The commercial precedent is genuine and worth stating precisely, because it is the strongest available evidence that self-similar geometry does real engineering work.

An antenna's resonant behaviour is set by its electrical length relative to wavelength. Two consequences of self-similarity follow:

1. Space-filling lowers the resonant frequency for a fixed physical envelope. A Koch monopole of a given height contains more conductor length than a straight monopole of the same height, so it resonates lower. For a handset — where the physical envelope is fixed by the industrial design and the frequency bands are fixed by regulators — this is the whole game.

2. Self-similarity at scale factor `s` produces resonances at frequency ratios approximately `s`. A structure containing copies of itself at half scale tends to resonate at `f`, `2f`, `4f`, `8f`. That is a multiband antenna obtained from *one* radiating element rather than from several, which is exactly what a device that must serve many bands in one envelope requires. Sierpinski and Koch monopoles were studied extensively for precisely this property from the 1990s onward, building on Nathan Cohen's earlier work.

Both of these are honest and load-bearing. Now the limit. The claim that fractality is *optimal* — that a fractal antenna necessarily outperforms a well-designed conventional multiband antenna of the same size — is contested in the antenna literature, and several careful comparisons have found conventional meander and inverted-F designs performing comparably or better. **The safe statement is: self-similarity is a reliable and elegant *route* to multiband, size-reduced behaviour, not a guarantee of superiority.**

How far the analogy carries into architecture It carries this far. A self-similar boundary offers usable features at several characteristic sizes simultaneously. A single geometry that presents 0.6 m recesses, 2.2 m alcoves, and 6.5 m bays serves the pair, the small group, and the assembly out of one construction — the same way a self-similar radiator serves several bands out of one element. That is a real economy and it is the reason to build this way.

It stops here. There is no scale-invariance principle for human occupancy corresponding to the scale-invariance of Maxwell's equations. People have one size. Hall's proxemic bands (Chapter 12 §12.4) are absolute distances, not ratios; Gehl's sensory thresholds are absolute distances. **A social geometry is self-similar over at most three or four scales before it runs into the fixed size of a human being at the bottom and the fixed limit of a sightline at the top.**

That is a bounded, honest, useful claim: **three to four scales, not infinitely many.** Which is exactly what §13.10 builds.

13.9 The crenellated boundary: more edge at no area cost

Theorem 13.5 in practice. Replace a straight boundary with a square-wave of tooth depth `d` and half-period `w`:

```
per period of length 2w, the path runs:  w  along the outer line
                                         d  inward
                                         w  along the inner line
                                         d  outward
```

```
path length per period = 2w + 2d
multiplier = (2w + 2d) / 2w = 1 + d/w
```

If the teeth alternate inward and outward about the original line, the enclosed area is unchanged exactly. Applied to a square room of 120 m^2 (side 10.954 m, perimeter 43.818 m):

d	w	multiplier	new perimeter
0.50 m	1.50 m	1.333	58.42 m
0.75 m	1.50 m	1.500	65.73 m
1.00 m	1.50 m	1.667	73.03 m
1.00 m	1.00 m	2.000	87.64 m
1.50 m	1.50 m	2.000	87.64 m

And it iterates. Applying the same `d/w = 2/3` ratio at each of three self-similar orders, with each order at one third the size of the previous:

order	d	w	multiplier this order	cumulative	perimeter
1	1.000 m	1.500 m	1.667	1.667	73.03 m
2	0.333 m	0.500 m	1.667	2.778	121.7 m
3	0.111 m	0.167 m	1.667	4.630	202.9 m

Order 3 is where the arithmetic detaches from reality and the honest engineer stops. A 111 mm recess at 167 mm pitch is not a seat, not a niche, and not reachable; it is a texture. It has real acoustic effects — at 111 mm scale it scatters sound above roughly 1.5 kHz, which is useful diffusion — but it does not carry a person, and pretending that perimeter metres at that scale count toward the Whyte ratio would be a straightforward fraud.

Rule: only count boundary at scales that can hold the thing you are counting for. For seating, that means `d >= 0.40 m` and `w >= 0.60 m`. For acoustic diffusion, much smaller scales count and Chapter 14 uses them. **This is the discipline that separates a fractal argument from fractal decoration.**

The inward-only case, which is the real one In a real building you rarely get to push teeth outward — there is a structural wall, a corridor, or another tenant out there. Inward-only crenellation converts open floor into niche floor:

For the same 120 m^2 room, recessing half the perimeter run (21.91 m) inward by `d`:

d	new perimeter	niche area created	remaining open floor	Whyte requirement on remaining floor	edge available
0.75 m	65.73 m	16.43 m^2	103.57 m^2	11.33 m	65.73 m
1.00 m	73.03 m	21.91 m^2	98.09 m^2	10.73 m	73.03 m
1.20 m	78.87 m	26.29 m^2	93.71 m^2	10.25 m	78.87 m

At `d = 1.00 m`: **the room gives up 21.9 m^2 of open floor — 18% — and gains 29.2 m of boundary, against a Whyte requirement that has simultaneously fallen to 10.73 m. It ends with roughly seven times the required sittable edge.**

And the 21.9 m^2 is not lost. It is exactly where the seating goes. **The floor area was converted into the thing the floor area existed to enable.** An open 120 m^2 room with 43.8 m of unusable perimeter is worse in every measurable respect than a 98 m^2 room with 73 m of occupied edge, and Chapter 12's worksheet will say so in numbers.

The construction need not be structural. Freestanding banquette backs, tall shelving, storage walls, planters and screens produce the same boundary multiplication, and they do it at a `delta` and `C` far below masonry.

13.10 The nested enclosure ladder

Now assemble the scales. Choose a scale ratio `r = 3` — chosen because it lands the rungs close to the human thresholds already established, not because three is magic:

k	Scale	`L_k = 0.75 * 3^k`	Area ~ L^2	What it holds	Threshold it matches
0	desk	0.75 m	0.6 m^2	one person, one task	Hall personal, 0.45-1.2 m
1	alcove	2.25 m	5.1 m^2	two to four	Hall social, 1.2-3.6 m
2	room	6.75 m	45.6 m^2	six to twelve	Gehl full-sensory, ~7 m
3	wing	20.25 m	410 m^2	thirty to eighty	between Gehl 7 m and 25 m
4	quarter	60.75 m	3691 m^2	a few hundred	Gehl "read emotion" ~25 m exceeded; scenery
5	campus	182.25 m	33215 m^2	thousands	Gehl "identify a person" ~100 m exceeded

Two of these rungs are doing something the others are not. **The desk-to-room span (k = 0 to 2) is the social ladder; above it, the rungs are logistical.** Past roughly 20 m, other people are not conversational partners and the enclosure's job changes from enabling contact to organising movement. §13.8's limit is visible here: **the self-similar social geometry is three rungs deep, maybe four. Claims beyond that are architecture, not sociology.**

For a two-dimensional plan, each rung holds `m = r^2 = 9` sub-units of the rung below, giving similarity dimension `D = ln 9 / ln 3 = 2.0000` — which is simply the statement that a plan fills a plane, as it should. **A plan whose measured `D` came out below 2 would be telling you it has gaps.** This is a sanity check, not a design target.

The important design consequence is the **multi-scale response in bounded volume**, which is Postulate 13.2's actual cash value. A room built with a crenellated boundary at `d = 1.0 m` and a secondary order at `d = 0.35 m` simultaneously offers:

- 0.35 m recesses: shoulder-scale, a place to stand out of the traffic, a shelf, a lean
- 1.0 m recesses: a seat with the back protected — Gehl's edge condition satisfied
- 2.25 m bays formed by pairs of recesses: a four-person alcove
- 6.75 m room: the whole assembly

One geometry, four occupancy scales, no moving parts, no power, and no booking system. That is what the fractal argument is for. It is not that the shape is beautiful; it is that a single construction serves four different group sizes and therefore satisfies Chapter 12's `N` requirement without four separate rooms.

13.11 The parameters that change with scale — hyperrecursive symmetry, measured

Postulate 13.1 says the same equation applies at every rung with different parameters. Here is the ladder with illustrative parameters. **These are illustrative values chosen to show the shape of the drift; they are not measurements, and any real project must fill this table from its own survey (Chapter 24).**

Scale	L (m)	`delta` (per pe- riod)	`tau` ~ 1/delta	`eta`	`N`	`ln N`	`eta` `ln N`	`Q`	`delta` `Q`	`C`	`h`
desk	0.75	0.600	1.7	0.40	3	1.0986	0.4394	0.2	0.1200	0.05	+0.2694
alcove	2.25	0.200	5.0	0.30	4	1.3863	0.4159	1.0	0.2000	0.12	+0.0959
room	6.75	0.080	12.5	0.25	6	1.7918	0.4479	3.0	0.2400	0.15	+0.0579
wing	20.25	0.030	33.3	0.15	10	2.3026	0.3454	20.0	0.6000	0.60	-0.8546
campus	60.75	0.010	100.0	0.08	16	2.7726	0.2218	150.0	1.5000	2.00	-3.2782

Read the columns as a set of monotone trends, because the trends are the content:

- **`delta` falls with scale.** A desk arrangement decays in days; a campus decays in decades. Larger structures are more durable, which is the point of building them.
- **`tau = 1/delta` rises correspondingly**, from under two periods to a hundred. Chapter 2's time constant: larger enclosures are slower filters. They respond to sustained forcing, not to impulses. **A one-day event changes a desk; it does not change a campus, and it is not a criticism of the event to say so.**
- **`eta` falls with scale.** Conversion of interaction into durable structure is most efficient in the smallest unit, where the person who had the idea is also the person who can act on it. Every rung upward inserts a handoff and every handoff loses.
- **`N` rises with scale, but only logarithmically in effect.** From desk to campus, `N` goes 3 to 16 and `ln N` goes 1.0986 to 2.7726 — a factor of 5.3 in categories buys a factor of 2.5 in ceiling. Chapter 8's brutal logarithm.

- **`eta ln N` is roughly flat and then falls.** The product of a rising `ln N` and a falling `eta` is nearly constant through the social rungs and declines above them.
- **`Q` explodes.** From 0.2 to 150 across the ladder.

That last row is the whole result:

step	`Q` grows by	`ln N` grows by
desk -> alcove	x 5.00	x 1.26
alcove -> room	x 3.00	x 1.29
room -> wing	x 6.67	x 1.29
wing -> campus	x 7.50	x 1.20

`Q` grows several-fold per rung while `ln N` grows by about a quarter. `delta Q` therefore outruns `eta ln N` and headroom goes negative as scale increases — in this illustration, somewhere between the room and the wing.

This is Chapter 6's "success consumes its own headroom" and Chapter 23's institutional ossification, arriving here through pure geometry. **The reason large organisations feel stuck is not that their people are worse. It is that `delta Q` scales with the stock and `eta ln N` does not.** The nested ladder makes it visible because the same equation is written five times with only the parameters changing, which is precisely what Postulate 13.1 claimed and precisely the form in which the claim can be checked.

The design consequence is direct and it is the reason this chapter sits in Part III: **push work down the ladder.** Whatever can be decided at the alcove should be, because `eta` is higher and `delta Q` is lower there. Chapter 21's rule — put the state in the medium, not in a coordinator — is the same instruction from the governance side.

13.12 Self-similar seating and the repeated pattern

Alexander's *A Pattern Language* is the acknowledged ancestor here (CANON §3). The observation this chapter adds is that **the same seating pattern, correctly stated, is valid at every rung, and that is what makes the geometry self-similar in a useful rather than a decorative sense.**

The pattern, stated once:

```
A person will occupy a position that has:
  (a) the back protected                    -- a wall, a high back, a recess
(b) a view out over a larger territory    -- toward the room, the bay, the court
   (c) a visible, low-cost exit               -- see Chapter 12 §12.5
(d) a boundary at reachable scale        -- something to set a cup on, lean on, touch
```

Now instantiate it at four scales:

Rung	(a) back protected by	(b) view out over	(c) exit	(d) reachable boundary
desk (0.75 m)	a 1.2 m screen or shell	the alcove	one step, unobstructed	desk edge, shelf lip

Rung	(a) back protected by	(b) view out over	(c) exit	(d) reachable boundary
alcove (2.25 m)	banquette back / storage wall	the room	alcove mouth, 1.2 m clear	ledge at 430 mm, table edge
room (6.75 m)	the crenellated wall	the bay or court	two doors, opposed	window ledge at 450 mm
wing (20.25 m)	a built edge, colonnade, arcade	the court	multiple, legible	parapet, planter, low wall

The pattern is invariant; only the dimensions scale. This is the strong form of hyperrecursive symmetry applied to furniture, and unlike Postulate 13.1's parameter claim it is directly falsifiable: build all four rungs, count occupancy at each, and see whether the pattern predicts where people go. Chapter 12's worksheet is the instrument.

The falsification condition, stated plainly: if positions satisfying (a)-(d) are not occupied preferentially over positions that fail one or more of them, the pattern is wrong for that population and this chapter's seating requirements should be discarded for that project. Say so in the survey. Do not defend the pattern against the count.

13.13 The modelling obligation

Every geometric claim in this chapter falls into one of three tiers, and the tier determines what you owe before you may build on it.

Tier 1 — Proven, use freely. Theorem 13.1 (focal reflection), Theorem 13.2 (focus-directrix), Corollary 13.3 (path equality), Theorem 13.4 (similarity dimension), Theorem 13.5 (crenellation multiplier). These are geometry. They are exact. They require nothing but correct arithmetic.

Tier 2 — Standard physics with stated approximations, use with the approximation declared. The `70 lambda/D` beamwidth relation, the `(pi D/lambda)^2 eta_ap` gain relation, the `D >= 2 lambda` frequency floor, Sabine's equation in Chapter 14. Each is an engineering approximation with a known domain of validity. **Declare `eta_ap`. Declare that near-field performance falls short. Declare that the figures are single-frequency.**

Tier 3 — Requires modelling or measurement of the specific geometry. May not be certified from the concept. Anything that depends on the detailed shape: the scattering behaviour of a particular crenellation, the diffusion coefficient of a particular relief, the acoustic sub-division actually achieved by a particular alcove depth, the occupancy response to a particular nested plan.

Tier 3 is where fractal claims live, all of them, without exception.

The acceptable methods for discharging a Tier 3 obligation, in ascending order of cost:

1. **Scale physical measurement.** A cardboard mock-up of an alcove and a phone measuring RT60 and a level difference at two positions is a real measurement and costs an afternoon.
2. **Geometrical-acoustics simulation (ray or image-source).** Valid above the Schroeder frequency, which for a 100 m^3 room at RT60 = 0.5 s is in the low hundreds of hertz. Cheap, fast, adequate for mid and high frequency.

3. **Wave-based simulation (FEM / BEM / FDTD).** Required below the Schroeder frequency and for any claim about resonance, focusing, or small-scale diffusion. Expensive.
4. **Full-scale prototype and measurement.** The only method that settles a Tier 3 question definitively.

A specification that invokes self-similarity without naming which of these four was used is not a specification. It is a mood board.

The Specification

R13.n are requirements. Each names its tier and the term of the law it acts on.

R13.1 (Tier 1; acts on C). Any spatial function that can be discharged by shape shall be discharged by shape rather than by a powered device, unless the powered device is shown to have a lower whole-life C. Show the comparison; do not assert it.

R13.2 (Tier 1; acts on S, u). Room boundaries intended to carry seating shall be crenellated to a computed multiplier of not less than 1.5, using `1 + d/w` with `d >= 0.40 m` and `w >= 0.60 m`. Boundary at scales below `d = 0.40 m` shall not be counted toward any seating requirement.

R13.3 (Tier 1; acts on N). A nested enclosure shall present usable features at not fewer than three distinct scales, with adjacent scales differing by a factor of not less than 2.5.

R13.4 (Tier 1; acts on N, u). The four-part seating pattern of §13.12 — back protected, view out, visible exit, reachable boundary — shall be satisfiable at every rung of the built ladder, and the surveyor shall record occupancy at each rung to test it.

R13.5 (Tier 2; acts on C_effort). Any acoustic reflector specified for speech shall have its frequency floor computed as `f_floor = 2c/D` and stated on the drawing. **A reflector whose floor lies above 2 kHz shall not be described as improving speech; it may be described as a consonant-band device.**

R13.6 (Tier 2). Any stated aperture gain shall carry its assumed `eta_ap`, shall be labelled an upper bound, and shall be accompanied by a note that near-field performance will fall short.

R13.7 (Tier 1; acts on C). Parabolic reflectors shall be specified in focal form `y = x^2/(4f)` with `f`, `D`, depth `D^2/(16f)` and `f/D` all stated. An `f/D` between 0.24 and 0.30 shall be preferred unless a reason is recorded.

R13.8 (Tier 2; acts on u). The focus of any listening reflector shall be positioned so that the intended listener's head occupies it, and the shell shall be aimed such that no other seat lies within the half-power beam. **A reflector that privileges more than one seat has created a hierarchy of seats, which Chapter 17 prohibits.**

R13.9 (Tier 3; acts on u, C_exclusion) — PROHIBITION. No concave surface with a radius of curvature comparable to a room dimension shall face the occupied zone unless it is (a) diffusive, (b) absorptive, or (c) focused demonstrably away from every seat, and unless the demonstration is by one of the four Tier 3 methods of §13.13.

R13.10 (Tier 3). Every claim in a specification that rests on self-similarity shall name the Tier 3 method used to verify it, the frequency or scale range over which the verification is valid, and the date. **A self-similarity claim without a named method shall be struck from the specification.**

R13.11 (Tier 3; acts on S). Acoustic sub-division of an open floor by alcove geometry shall be verified by measuring the A-weighted level difference between the alcove seat and a reference position at 4 m, with a real source. A design target of not less than 6 dB shall be set; the measured value shall be recorded whether or not it meets the target.

R13.12 (Tier 1; acts on `delta`). Geometric elements intended to perform a function shall be constructed in materials whose service life exceeds the intended function's horizon by a factor of two. The `delta` advantage of geometry is entirely a material property and is forfeit if the geometry is built in something that fails first.

R13.13 (acts on `C`, and on dignity). Every recess, alcove and niche shall be cleanable, reachable, lit, and visible from at least one circulation route. **An unobservable recess is a liability, not an amenity**, and a geometry that creates concealed space without sightline creates a safety problem that no acoustic benefit offsets.

R13.14 (ethics, CANON §5). No nested geometry shall be arranged such that an occupant of an inner scale cannot leave without passing through an occupied outer scale. Nesting must not become entrapment. Every rung shall have an exit that does not require crossing another group.

R13.15 (measurement integrity). Perimeter counted toward the Whyte ratio shall be counted at seating scale only, per R13.2. Any published perimeter figure shall state the scale at which it was measured, because — by Theorem 13.4 — a boundary length is meaningless without its ruler.

R13.15 is worth dwelling on. **The length of a fractal boundary is not a number; it is a function of the measuring scale.** A specification that reports "78 m of edge" without saying at what scale has reported nothing. This is the Koch curve's lesson applied as a documentation rule, and it is the most common way fractal architecture claims become unfalsifiable.

The Failure Modes

F13.1 — Fractal Decoration

Damages: the credibility of every other number in the specification. **Description.** Self-similar ornament applied to a surface, described as if it delivered performance. Recesses too small to hold anything, relief too shallow to scatter at the frequencies that matter, a dimension quoted with no measurement behind it. **Signature.** The word "fractal" appears in the design narrative and no Tier 3 method appears anywhere. Perimeter is quoted without a scale. Occupancy at the fine scale is zero. **Detects with.** R13.10 and R13.15. Ask which of the four methods was used. If the answer is a rendering, the claim is decoration. **Fix.** Either commission the model, or delete the claim and keep the ornament as ornament. There is nothing wrong with ornament; there is a great deal wrong with ornament billed as engineering.

F13.2 — The Concave Room

Damages: `u`, `C_exclusion`. **Description.** Domed, vaulted, or curved surfaces facing the occupied zone, concentrating sound at focal points and producing seats that are acoustically privileged and seats that are acoustically dead. **Signature.** Occupants report that "you can't hear anything from that end" and "she's always too loud." A level sweep with a fixed source shows variation exceeding 6 dB between seats at equal distance. Discrete echo audible on a hand clap at specific

positions. **Detects with.** A hand clap and a walk. Then a level sweep with a calibrated source. R13.9 should have prevented it. **Fix.** Diffuse or absorb the concave surface. Do not attempt to "balance" it with electronics — that adds a permanent `C` line to fix a geometry error and it fails whenever the power does.

F13.3 — The Frequency Floor Ignored

Damages: trust; wasted capital. **Description.** A small reflector, screen, shell or "acoustic pod" specified to control speech, whose aperture is far smaller than the wavelengths in question. Sound diffracts around it and the device does approximately nothing at the frequencies that carry vowel energy. **Signature.** Measured level difference across the device is under 3 dB below 1 kHz. Occupants report that the pod "doesn't really work" while the specification claims 12 dB. **Detects with.** `f_floor = 2c/D`, computed in ten seconds. R13.5. **Fix.** Enlarge, or re-describe honestly as a consonant-band and privacy device, or replace with absorption, which has no frequency floor of this kind (though it has its own thickness-related low-frequency limit — see Chapter 14).

F13.4 — The Unreachable Rung

Damages: N, and the honesty of the category audit. **Description.** A nested plan whose fine scale is not actually usable — recesses that a person cannot sit in, alcoves whose entry is blocked, a "quiet corner" with no seat. The plan claims four occupancy scales and delivers two. **Signature.** Chapter 12's category audit strikes rows that the drawings claimed. Time-lapse shows the fine scale unoccupied at all times. **Detects with.** Worksheet B, plus one week of counting. **Fix.** R13.2's dimensional floor, applied before construction rather than after.

F13.5 — The Concealed Recess

Damages: safety, and therefore everything. **Description.** Boundary multiplication produces recesses with no sightline from any circulation route. In a workplace this generates avoidance; in public space it generates worse. **Signature.** Occupants — reliably, a non-random subset of them — avoid a whole zone. The zone appears clean in the count sheets because nobody is in it. **Detects with.** Sightline analysis on the plan, plus asking. This is one of the two failure modes where the intercept instrument is the correct one. **Fix.** R13.13. Every recess visible from a route. Jacobs' eyes on the street (CANON §3), at room scale. **Note.** A count sheet cannot distinguish "unused because unattractive" from "unused because avoided". **Do not diagnose from the count alone.**

F13.6 — Scale Inflation

Damages: the model's credibility. **Description.** Applying Postulate 13.1 past the point where its parameters were measured — claiming that because the desk and the room obey the same recursion, the nation does too, with parameters assumed rather than surveyed. **Signature.** A parameter table with no measurement dates and no methods. `delta` and `eta` for a city, asserted. **Detects with.** Chapter 24's standard: every parameter carries a date, a method, and an observer. **Fix.** Measure or stop. **The scale-free form of the law is a licence to model at a new scale; it is not a licence to skip the measurement at that scale.**

F13.7 — The Passive Device With Nothing To Concentrate

Damages: nothing directly; wastes capital and misdirects diagnosis. **Description.** Denge's actual lesson, misapplied. Excellent geometry installed in a system whose S is near zero — a beautifully sub-divided floor for a group with `u = 0`, or a listening niche in an organisation where nobody has anything to say that they are permitted to say. **Signature.** All geometric requirements pass; S and `eta` are unchanged; `h` does not move. **Detects with.** Chapter 12's worksheet section E, before and after. **Fix.** Chapters 9 and 11 own this. Geometry sets coefficients; it does not supply the input. **Stop building and go and find out why there is no signal.**

Worked Example

The subject

The same 120 m^2 departmental common room surveyed in Chapter 12. Its diagnosis was: N = 2, Whyte ratio 0.050 against a requirement of 0.109, OAR 0.185, `delta = 0.0808/week`, `eta_hat = 0.0455`, `h = −1.396`. The prescription from Chapter 12 was to raise N from 2 to 5, cut the acoustic `C`, and shed `Q`. **This example designs the geometry that delivers the N increase, and prices what it can and cannot certify.**

Constraints: no structural walls may be added (tenant fit-out only); the two glazed walls must remain glazed; the budget is fit-out joinery, not construction.

Step 1 — Boundary multiplication

Existing: 120 m^2, approximately square at 10.95 m x 10.95 m, perimeter 43.82 m. Of that, 21.9 m is glazing (which cannot be recessed) and 21.9 m is solid wall plus the kitchenette run.

Crenellate the 21.9 m of solid boundary inward, freestanding, using banquette backs and storage joinery. Choose `d = 1.00 m`, `w = 1.50 m`:

```
multiplier          = 1 + d/w = 1 + 1.00/1.50 = 1.6667
perimeter           = 43.82 -> 73.03 m
boundary gained     = 29.21 m
niche floor created = 21.91 m of run x 1.00 m = 21.91 m^2
open floor          = 120 - 21.91 = 98.09 m^2
```

Check against the Whyte ratio, recomputed on the reduced open floor:

```
required sittable edge = 98.09 x 0.1094 = 10.73 m
available boundary at seating scale (d >= 0.40 m) = 73.03 m
surplus factor = 6.8x
```

The requirement was never hard to meet. It was hard to meet in a room with no boundary. Fitting seats to the new edge at 0.65 m pitch would allow 112 positions, which is absurd — the room has 90 users. The correct move is to occupy a fraction of the new boundary generously rather than all of it meanly.

Specified occupation of the 73.03 m:

Element	Boundary consumed	Seats	Rung
Four banquette alcoves, `d = 1.00 m`, 2.4 m wide	4 x 4.4 = 17.6 m	16	alcove (k=1)
Deepened window ledges, 450 mm, perch height 700 mm	9.0 m	9 perch	desk (k=0)
High table + rail against the kitchenette crenellation	3.6 m	6 stool	desk/alcove
Existing free tables, unchanged, in the open floor	0	24	room (k=2)
Unoccupied boundary (circulation, storage face, doors)	42.8 m	0	–
Total	**73.0 m**	**55**	

Step 2 — The scale ladder, checked against R13.3

Rung	Built as	Characteristic dimension	Ratio to rung below
k=0 desk	Ledge perch, high-table stool	0.75 m	–
k=1 alcove	Banquette recess, 2.4 m wide x 1.0 m deep	2.4 m	3.20
k=2 room	The 98 m^2 open floor with its tables	9.9 m	4.13

Three rungs, adjacent ratios 3.20 and 4.13, both above the 2.5 minimum. **R13.3 satisfied.**

Secondary order: the banquette backs carry a 0.30 m deep x 0.45 m wide relief for shelving and personal items. This is below R13.2's `d >= 0.40 m` seating floor, so **it contributes zero to the Whyte ratio and is not counted.** It is specified for a different reason (§13.9's acoustic scattering, and R13.12's reachable-boundary requirement (d) from the seating pattern), and that reason is stated on the drawing.

Step 3 — Category audit, recomputed (Chapter 12 §12.10 procedure)

#	Posture	Purpose	Condition	Available?
1	Sitting upright, standard chair	Eating / pair talk	Bright, live, open	Y

#	Posture	Purpose	Condition	Available?
2	Sitting, backed banquette	Small-group deliberation	Dimmable, quieter, enclosed	Y — new
3	Perching, 700 mm ledge	Brief stop, solo, daylight	Bright, live, edge	Y — new
4	Standing at rail	Pair work, transit-adjacent	Bright, live	Y — new
5	Standing, kettle queue	Waiting	Bright, live	Y
6	Sitting reclined	Solo rest	Dim	N — not provided

Surviving distinct categories: **5**.

```
N: 2 -> 5
ln N: 0.6931 -> 1.6094      (x 2.32)
```

This is the increase Chapter 12's intervention plan required, and it was delivered entirely by boundary geometry plus joinery. **No wall was built. No square metre left the lease.**

Step 4 — The parabolic element

One listening niche is specified, at the quiet end of the room, as a demonstration piece and as a functional device for a single person taking a call or listening to the room.

Geometry:

```
D    = 1.2 m     (aperture, set by the banquette module width less returns)
 f     = 0.35 m
 y     = x^2 / (4 * 0.35) = x^2 / 1.40
 depth = D^2/(16f) = 1.44 / 5.60 = 0.2571 m
f/D   = 0.2917                       -- within R13.7's preferred 0.24–0.30
 rim   at x = 0.60 m, y = 0.2571 m
 |PF| at rim = y_rim + f = 0.6071 m      -- Theorem 13.2, checked
focus at 350 mm from the vertex, i.e. at seated ear height when the seat is set correctly
```

Performance, computed and labelled per R13.5 and R13.6:

```
f_floor = 2c/D = 2 x 343 / 1.2 = 572 Hz     <-- STATED ON THE DRAWING
```

Frequency	Beamwidth	Aperture gain, `eta_ap = 0.5`, **UPPER BOUND**
250 Hz	80.0 deg	5.8 dB — below floor, disregard
500 Hz	40.0 deg	11.8 dB — at floor, treat as marginal
1000 Hz	20.0 deg	17.8 dB
2000 Hz	10.0 deg	23.8 dB
4000 Hz	5.0 deg	29.9 dB

Frequency	Beamwidth	Aperture gain, `eta_ap = 0.5`, **UPPER BOUND**

The drawing note, written out in full, because this is what R13.6 is for:

> Acoustic niche NC-01. Parabolic, `y = x^2/1.40`, aperture 1200 mm, focal length 350 mm, depth 257 mm. Useful-directivity floor 572 Hz (`D = 2 lambda`, c = 343 m/s at 20 C). Gain figures are theoretical aperture bounds at `eta_ap = 0.5` and will not be realised: the intended source is in the near field, and the listener's head partially blocks the aperture. Expected realised benefit is a consonant-band (2-8 kHz) improvement in one direction. **This element does not amplify speech.** Verification: Tier 3 method (1), physical mock-up, level difference measured at the focus versus 1 m off-axis, both with a 1 kHz-centred pink-noise source at 2 m on-axis. Record and attach.

Aimed per R13.8 so that its half-power beam at 1 kHz (20 degrees) contains no other seat: the niche faces the room diagonally across the open floor, and the nearest other seat lies 34 degrees off axis.

And per R13.9, the reverse check: the niche's concave surface is 1.2 m across in a room whose smallest dimension is 2.8 m. That is not "comparable to a room dimension" at the 1:2 threshold this manual uses, and its focus is 350 mm from its own vertex — inside the niche, not out in the room. The prohibition does not bite. Record the check; do not assume it.

Step 5 — What is certified and what is not

Per §13.13, the specification's claims sorted by tier:

Claim	Tier	Status
Perimeter multiplies by 1.6667	1	**Certified.** `1 + d/w`, exact.
Boundary rises 43.82 -> 73.03 m	1	**Certified.** Arithmetic.
Niche floor 21.91 m^2, open floor 98.09 m^2	1	**Certified.**
Whyte requirement 10.73 m, met 6.8x over	1	**Certified** (at seating scale, per R13.15).
Three rungs at ratios 3.20 and 4.13	1	**Certified.**
`N` rises 2 -> 5, `ln N` 0.693 -> 1.609	1 (audit)	**Certified by procedure**, subject to the strike rule being applied honestly.
Parabola focuses at 350 mm	1	**Certified.** Theorem 13.1.
Niche frequency floor 572 Hz	2	**Declared with approximation** (`D = 2 lambda` rule).
Niche gain figures	2	**Upper bounds only.** Not certified. Mock-up required.

Claim	Tier	Status
Alcoves achieve >= 6 dB level difference from the open floor	**3**	**NOT CERTIFIED.** R13.11 requires measurement. Budget one afternoon and a mock-up.
The 0.30 m secondary relief scatters usefully above ~1.5 kHz	**3**	**NOT CERTIFIED.** Plausible from the scale, unproven for this shape. Either model it or delete the claim.
Occupancy will follow the four-part seating pattern at all three rungs	**3**	**NOT CERTIFIED.** Falsifiable by Chapter 12's count. Predicted, not promised.

Three of the twelve claims are uncertified, and they are the three that sound most impressive. That ratio is normal and it is the reason §13.13 exists.

Step 6 — The Tier 3 discharge, actually performed

Two afternoons of work, method (1), physical measurement on a mock-up built from 18 mm ply and 50 mm mineral wool.

Test A — alcove sub-division (R13.11). Source: a loudspeaker at the centre of the open floor reproducing pink noise at 70 dBA at 1 m. Measurement: A-weighted level at the alcove seat position versus at a reference 4 m from the source in the open.

```
reference position, 4 m open floor        : 58.3 dBA
alcove seat, same distance, inside recess : 53.6 dBA
level difference                          : 4.7 dB
design target                             : 6.0 dB
RESULT: TARGET NOT MET
```

Recorded as not met, per R13.11's requirement to record the measured value whether or not it meets the target. The deficiency is diagnosed as insufficient recess depth relative to wavelength: at `d = 1.00 m`, the shielding is effective only above roughly `c/(2d) = 172 Hz` for path-length reasons but is compromised by the open top of the banquette back, which allows sound over the top at all frequencies.

Remedy, at negligible cost: raise the banquette back from 1.35 m to 1.75 m and line its room-side face with 50 mm absorption. Re-measured:

```
alcove seat, back raised to 1.75 m, lined : 51.4 dBA
level difference                          : 6.9 dB
RESULT: TARGET MET
```

This is what a Tier 3 obligation costs and what it is worth. The first design failed by 1.3 dB. Nothing in the fractal argument would have revealed that, because the fractal argument is about boundary length and this is about a wavelength and an open top. A specification that had certified the alcove from the concept would have shipped a room that did not sub-divide.

Test B — the parabolic niche. Source at 2 m on-axis, pink noise. Level at the focus versus 1 m off-axis, in third-octave bands:

```
250 Hz band : +0.4 dB  (below floor, as predicted; the shell is invisible at this wavelen
 500 Hz band : +2.1 dB
 1000 Hz band: +4.8 dB
 2000 Hz band: +6.3 dB
 4000 Hz band: +7.1 dB
```

Compare with the theoretical bounds: 17.8 dB at 1 kHz, realised 4.8 dB. **The realised figure is 27% of the theoretical bound in decibels, and this is entirely expected** — near-field source, head blockage, finite surface accuracy, edge diffraction. The device does what §13.4 said it would do: a consonant-band directivity improvement of a few decibels, and nothing below 500 Hz.

Report it that way. A 4.8 dB improvement at 1 kHz rising to 7.1 dB at 4 kHz is a genuinely useful device and a completely honest description. "17.8 dB of acoustic gain" would have been a lie that the arithmetic in this chapter itself would have supported.

Step 7 — What the geometry did to the law

Parameter	Before	After the geometry	Mechanism
`N`	2	5	Boundary multiplication -> four new posture/purpose/condition categories
`ln N`	0.6931	1.6094	–
Sittable edge	6.0 m	55 seats on 30.2 m of occupied boundary	Crenellation
Whyte ratio	0.050	0.308 (occupied edge / open floor)	–
`C_effort` (acoustic)	high	reduced; alcoves 6.9 dB quieter than open floor	R13.11 verified
`delta`	0.0808/wk	unchanged by geometry	Geometry does not touch this; Chapter 12's register does
`eta`	0.0455	predicted to rise; **not** claimed	Requires Chapter 12's artefact discipline
`S`	22 contacts/wk	predicted to rise; **not** claimed	Geometry sets coefficients, not input (F13.7)

The geometry moved `N` and `C`. It did not move `delta`, and it cannot be credited with moving `eta` or `S` until those are measured. This is the discipline the chapter has been arguing for, and it is the reason the table has "not claimed" in it twice.

Recomputing headroom with only the certified changes and the measured `eta` held at its old value:

```
eta ln N = 0.0455 x 1.6094 = 0.0732
delta Q  = 0.0808 x 9      = 0.7272        (Q not yet shed)
C        = 0.70 - 0.55 (acoustic) = 0.15
h        = 0.0732 - 0.7272 - 0.15 = -0.8040
```

Improved from −1.396 to −0.804 on geometry alone, which is 42% of the gap. The remaining moves are Chapter 12's — shed Q, install the register — and Chapter 14's full acoustic treatment.

Geometry did what geometry does: it raised the ceiling and cut the cost, permanently, with no power and no ongoing administration, and it did not pretend to do more.

What this example demonstrates

1. **Theorems 13.4 and 13.5 turned an unmeetable requirement into a 6.8x surplus with joinery**, at the cost of 18% of the open floor which became the seating it was needed for.
2. **The three-rung ladder delivered N = 5 from one construction**, which is Postulate 13.2's practical payoff: multi-scale response in bounded volume.
3. **The parabola performed exactly as the arithmetic predicted, at roughly a quarter of the theoretical bound in decibels**, and the honest description of it is more useful than the flattering one.
4. **The one claim that required Tier 3 verification failed on first build and was fixed for the price of 0.4 m of plywood.** Had it been certified from the concept, the room would have shipped with alcoves that did not work, and nobody would have known why.

Chapter 14 takes the acoustic term that appeared in every cost line of this example and computes it completely.

Chapter 14 . Acoustics for Turn-Taking

This is the most directly actionable chapter in the manual. Everything in it can be measured with equipment costing less than a chair, computed with arithmetic you can do on paper, and fixed with materials available from any building supplier.

It is also the chapter where the manual's central sociological claim touches physics hard enough to make a sound.

Woolley, Chabris, Pentland, Hashmi and Malone (*Science*, 2010) found a single statistical factor c predicting a group's performance across varied tasks, not predicted by members' average or maximum IQ, and correlating with three things: social sensitivity, **equality of conversational turn-taking**, and the proportion of women. Chapter 18 handles c in full. This chapter takes one of the three correlates and observes something that the sociology literature does not usually say:

Turn-taking is substantially a property of the room.

Not entirely. Status, culture, chairing, agenda and personality all bear on who speaks. But a measurable and often dominant share of turn-taking inequality in a physical meeting is produced by acoustics — by reverberation time, by background level, by the direct-to-reverberant ratio at each seat, and by the Lombard feedback loop those three set up between the people in the room.

The consequence is the chapter's thesis:

Acoustic design is c-factor engineering, and it is the cheapest c-factor intervention available.

And the corollary, which is the sentence to remember:

A hard-surfaced reverberant room silences the quiet voice by physics, not by dominance. Nobody in the room is behaving badly. The room is doing it, continuously, to everyone, and it will keep doing it after every training course, every facilitation technique, and every exhortation to "make space for others."

The Law

What is proven, and what "proven" means here

Sabine's equation — EMPIRICAL LAW WITH A DERIVATION, not a theorem of pure mathematics.

```
RT60 = 0.161 * V / A          V in cubic metres
                  A in metric sabins (square metres of equivalent absorption)
                              RT60 in seconds
```

RT60 is the time for the sound energy density in the room to fall by 60 dB after the source stops. A is the total absorption:

```
A = sum over surfaces of ( S_i * alpha_i )  +  sum over objects of ( A_j )

    S_i     area of surface i, m^2
  alpha_i absorption coefficient of surface i at the band of interest, 0 to 1
```

```
A_j    absorption of a discrete object (a person, a chair), in m^2 sabins each
```

The constant 0.161 carries units of s/m and is derived from the speed of sound and the mean free path between reflections in a rectangular enclosure. **The equation is not a curve fit; it follows from a statistical-energy argument. But the argument makes assumptions, and when the assumptions fail the equation fails, sometimes badly. §14.2 lists the assumptions and §14.4 shows a case where Sabine over-predicts by 78%.**

The inverse-square law and the two-field model — PROVEN, standard acoustics.

```
L_p(r) = L_W + 10 log10( Q/(4 pi r^2)  +  4/R )

    first term  = the DIRECT field, falls 6 dB per doubling of distance
    second term = the REVERBERANT field, the same everywhere in the room
    Q           = source directivity factor (a talker on a table, Q ~ 2)
    R           = A / (1 - alpha_bar)   the "room constant", m^2
    alpha_bar   = A / S_total
```

Critical distance — PROVEN, from setting the two terms equal.

```
d_c = 0.141 * sqrt( Q * R )      metres

Beyond d_c the reverberant field dominates and moving closer to the talker stops helping.
```

Decay rate — PROVEN, by definition of RT60. Sound decays at `60/RT60` decibels per second, so a gap of `t` seconds leaves the previous talker's energy `60 t / RT60` decibels down.

What is cited (CANON §3)

Woolley et al., *Science* 2010. The c factor. Correlates: social sensitivity, equality of conversational turn-taking, proportion of women. **Honest limit, per the CANON and per the authors: c was measured in short-task laboratory settings. Extrapolation to a standing institutional committee is a modelling assumption, not a finding.**

ANSI/ASA S12.60 sets, for classrooms, an unoccupied background noise limit of 35 dBA and a reverberation time limit of 0.6 s for rooms up to 283 m^3 (0.7 s for 283-566 m^3). This is a real published standard, arrived at through the speech-intelligibility literature, and it is the best available anchor for deliberation-room targets. **It is a classroom standard. Applying it to a meeting room is a transfer, and the transfer is defensible — a committee member straining to hear is in the same acoustic situation as a child at the back — but it is a transfer, and it should be labelled one.**

What is postulated

POSTULATE 14.1 (badged): a measurable and often dominant share of turn-taking inequality in a physical meeting is caused by the room's acoustics. The mechanism is given in §14.6 and §14.7 and is entirely physical; the *share* is not established in the literature and must be measured per-room. **This manual does not claim a number for it. It claims a mechanism, computes the mechanism's magnitude, and tells you how to measure whether it dominates in your room.**

POSTULATE 14.2 (badged): therefore acoustic treatment is a c-factor intervention. This chains Postulate 14.1 to Woolley's correlation. Note carefully what the chain requires: it requires

that turn-taking equality *causes* c rather than merely correlating with it. **Woolley et al. report a correlation. The causal direction is not established.** If turn-taking equality is a symptom of c rather than a driver, acoustic treatment will raise turn-taking equality and leave c unchanged. **§14.13 gives the falsification test, and any honest programme must run it.**

POSTULATE 14.3 (badged): acoustically-caused silence is indistinguishable, in the room's own records, from voluntary silence. A contribution never attempted leaves no trace. This is why §14.9 computes exclusion from physics rather than from self-report, and why Chapter 12's `C_exclusion` term could not be measured any other way.

The honesty clauses

- **`ln N` is a hard ceiling per interaction** (Paper D, Cor 2.2). When §14.9 computes the loss from silencing three of eight voices, it is computing a loss of *ceiling*, not of realised surplus.
- **As deployed, the surplus functional is a LOWER BOUND** on the ensemble's information.
- **Paper A makes no empirical or metaphysical claim.** The bridge from `ln N` to a real meeting is interpretation, and it is badged.
- **Every absorption coefficient in this chapter is NOMINAL.** They are typical published values for generic material descriptions. **They must be verified against the specific manufacturer's tested data for the specific product, tested to a recognised standard, before they are used in a specification that someone will build.** A "50 mm mineral wool panel" spans a range of real performance wide enough to change a design decision.

The Engineering

14.1 The chain of inference, stated as a chain

Set it out so each link can be attacked separately.

```
     1. Speech intelligibility at a listener depends on the direct-to-
reverberant ratio
     and on the signal-to-noise ratio at that listener.     [PROVEN PHYSICS]

   2. Both of those degrade with distance, with reverberation time,
      and with background level.                            [PROVEN PHYSICS]

  3. A quiet talker produces less sound power than a loud one, so at any given
      seat the quiet talker's D/R and SNR are worse.        [PROVEN PHYSICS]

   4. A talker who is not understood is interrupted, asked to repeat, or
    gives up, and each of those reduces their share of turns.  [OBSERVABLE; MECHANISM]

   5. Therefore the room produces turn-taking inequality that is not
      attributable to anyone's behaviour.                   [POSTULATE 14.1]

   6. Turn-taking equality correlates with the collective intelligence
     factor c.                                        [WOOLLEY 2010, CITED]
```

```
   7. Therefore acoustic treatment raises c.                [POSTULATE 14.2 -
- requires
                                                    causal direction in step 6]
```

Links 1 to 3 are physics and are not in dispute. Link 4 is a mechanism this manual asserts and tells you how to measure. Links 5 and 7 are badged postulates. Link 6 is a citation with a stated limit.

That is the whole argument, with its joints exposed. Anyone who wants to reject the conclusion must say which link they are cutting.

14.2 Sabine's equation: assumptions, and where it breaks

Sabine's derivation assumes a **diffuse field**: sound energy uniformly distributed throughout the volume, with all propagation directions equally probable, and absorption distributed reasonably evenly over the boundary. From that assumption, the mean free path between reflections in a rectangular room is `4V/S`, and the energy decay follows.

Five conditions must roughly hold. When they do not, say so.

Condition 1 — the room must be above its Schroeder frequency. Below it, the room's behaviour is dominated by discrete resonant modes, not by a statistical field, and RT60 is not a meaningful single number.

```
f_s = 2000 * sqrt( RT60 / V )      Hz
```

V	RT60	f_s
100 m^3	0.40 s	126.5 Hz
100 m^3	0.50 s	141.4 Hz
100 m^3	1.00 s	200.0 Hz
100 m^3	2.00 s	282.8 Hz
336 m^3	0.70 s	91.3 Hz
500 m^3	0.80 s	80.0 Hz
2000 m^3	1.20 s	49.0 Hz
10000 m^3	1.80 s	26.8 Hz

In a 100 m^3 meeting room, the Schroeder frequency sits around 130-140 Hz. That means the 125 Hz octave band — the leftmost column of every table in this chapter — is at or below the limit of Sabine's validity and must be treated as indicative only. For an 8.0 x 5.0 x 2.5 m room the axial modes are at 21.4, 34.3, 42.9, 64.3, 68.6 (two coincident), 85.8, 102.9, 137.2 (two coincident) Hz — the coincidences at 68.6 and 137.2 Hz are exactly the kind of modal pile-up that produces a boomy room, and no RT60 figure will reveal them.

Condition 2 — absorption must be reasonably low. Sabine over-predicts `RT60` as `alpha_bar` rises, because the equation permits `RT60` to remain finite as `alpha_bar` approaches 1, which is physically wrong (a room with total absorption has no reverberation at all). The Eyring correction fixes the limit:

```
 Sabine:  RT60 = 0.161 V / A
Eyring:  RT60 = 0.161 V / ( -S * ln(1 - alpha_bar) )     alpha_bar = A / S
```

Computed for the four build states of §14.4, in a 100 m^3 room of 145 m^2 total surface:

State	Band	A	alpha_bar	Sabine RT60	Eyring RT60	Sabine over-predicts by
A hard shell	500 Hz	8.73	0.060	1.844 s	1.788 s	3.1%
A hard shell	1000 Hz	8.24	0.057	1.954 s	1.898 s	3.0%
B office fit	500 Hz	35.53	0.245	0.453 s	0.395 s	14.7%
B office fit	2000 Hz	61.63	0.425	0.261 s	0.201 s	30.2%
C specified	500 Hz	36.94	0.255	0.436 s	0.378 s	15.4%
C specified	4000 Hz	41.52	0.286	0.388 s	0.329 s	17.8%
D over-treated	1000 Hz	94.52	0.652	0.170 s	0.105 s	**61.9%**
D over-treated	2000 Hz	104.99	0.724	0.153 s	0.086 s	**77.8%**

Rule: use Sabine below `alpha_bar` = 0.25, use Eyring above it, and state which you used. In the dead-room case the two formulae disagree by nearly a factor of two, and a specification that quoted Sabine without saying so would be describing a room 78% more reverberant than the one that got built.

Condition 3 — absorption must be reasonably evenly distributed. A room with all its absorption on one surface (the classic: an absorptive ceiling and five hard surfaces) does not develop a diffuse field. Sound bouncing in the horizontal plane never meets the ceiling and decays far more slowly than Sabine predicts. **This is the single most common real-world failure of the equation and it is why §14.14's specification requires absorption on at least two non-parallel planes.**

Condition 4 — the room must not be strongly coupled to another volume. An open door to a corridor, a folding partition, a double-height void: these create coupled spaces with two-slope decay curves that no single RT60 describes.

Condition 5 — dimensions must not be wildly disparate. Corridors, and rooms far longer than they are wide or tall, violate the mean-free-path derivation.

Air absorption adds a `4mV` term to `A` and matters above about 2 kHz in large volumes. In a 100 m^3 room at ordinary humidity it is negligible and is omitted throughout this chapter. In a 2000 m^3 hall it is not, and omitting it there would be an error.

14.3 Absorption coefficients

Every value in this table is NOMINAL. These are typical published values for generic material descriptions, given per octave band. They are adequate for the design arithmetic that follows and for deciding what to build. **They are not adequate for a construction specification. Before you buy anything, obtain the manufacturer's data for the specific product, tested to a recognised standard, at the specific mounting condition you intend — mounting changes low-frequency performance dramatically, and a panel tested with an air gap behind it will not perform the same way glued flat to a wall.**

Material (nominal)	125	250	500	1k	2k	4k
Concrete, sealed or painted; terrazzo	0.01	0.01	0.02	0.02	0.02	0.03
Concrete block, painted	0.10	0.05	0.06	0.07	0.09	0.08
Concrete block, coarse / unpainted	0.36	0.44	0.31	0.29	0.39	0.25
Gypsum board 13 mm on studs, air cavity	0.29	0.10	0.05	0.04	0.07	0.09
Plaster on solid backing	0.02	0.02	0.03	0.04	0.05	0.05
Glass, heavy plate	0.18	0.06	0.04	0.03	0.02	0.02
Glass, ordinary window	0.35	0.25	0.18	0.12	0.07	0.04
Wood floor on joists	0.15	0.11	0.10	0.07	0.06	0.07
Plywood panel 10 mm over air space	0.28	0.22	0.17	0.09	0.10	0.11
Carpet, heavy, on concrete	0.02	0.06	0.14	0.37	0.60	0.65
Carpet, heavy, on resilient underlay	0.08	0.24	0.57	0.69	0.71	0.73
Heavy velour drape, draped to half its flat area	0.14	0.35	0.55	0.72	0.70	0.65

Material (nominal)	125	250	500	1k	2k	4k
Mineral-fibre ceiling tile, suspended with plenum	0.20	0.40	0.60	0.70	0.75	0.75
Mineral wool panel 25 mm, direct to wall	0.10	0.30	0.65	0.85	0.90	0.90
Mineral wool panel 50 mm, direct to wall	0.20	0.55	0.85	0.95	0.95	0.90
Mineral wool panel 100 mm with 50 mm air gap	0.45	0.85	0.95	0.98	0.95	0.90
Bookcase, filled, open front	0.25	0.30	0.30	0.30	0.30	0.30
Water surface	0.008	0.008	0.008	0.008	0.008	0.008

Discrete absorbers, in m^2 sabins **each**, not per m^2:

Object (nominal)	125	250	500	1k	2k	4k
Adult, seated, on a lightly upholstered chair	0.15	0.25	0.35	0.42	0.46	0.46
Empty lightly upholstered chair	0.10	0.15	0.18	0.20	0.22	0.22

Four things to read out of this table before using it.

Porous absorbers fail at low frequency, and thickness is the reason. A porous absorber works by viscous loss as air moves through it, and air motion is greatest a quarter-wavelength from a rigid boundary. At 125 Hz a quarter wavelength is 0.69 m. **A 25 mm panel simply cannot reach the low-frequency air motion; a 100 mm panel with a 50 mm gap gets much closer, which is why its 125 Hz figure is 0.45 against the 25 mm panel's 0.10.** Every low-frequency problem in this chapter traces to this one fact.

Gypsum board on studs is a low-frequency absorber. Its 0.29 at 125 Hz is not a typo: a panel on a cavity is a mass-spring system that resonates and dissipates in the bass. It is the only absorption most ordinary rooms have down there.

People are significant absorbers and their absorption is mid-and-high biased. Eight seated adults contribute 2.8 m^2 sabins at 500 Hz and 3.4 at 1 kHz. In a small hard room that is a material fraction of the total, which is why **the occupied and unoccupied RT60 of a small room differ substantially and every specification must say which one it means.**

Carpet on concrete does almost nothing below 500 Hz. Its 125 Hz figure of 0.02 is the same as bare concrete. Carpet is a treble absorber, and a room "treated" with carpet alone gets darker, not clearer.

14.4 The 100 m^3 room, computed

Room: 8.0 m x 5.0 m x 2.5 m = 100.0 m^3 exactly. Floor 40 m^2, ceiling 40 m^2, walls `2(8 x 2.5) + 2(5 x 2.5)` = 65 m^2, of which 12 m^2 is glazing. Total surface `S` = 145 m^2. Eight occupants.

First, the shape of the relationship Before any material, see what `A` has to be. Sabine, `V = 100`:

A (m^2 sabins)	RT60
4	4.025 s
6	2.683 s
8	2.013 s
10	1.610 s
12	1.342 s
15	1.073 s
20	0.805 s
25	0.644 s
30	0.537 s
35	0.460 s
40	0.403 s
50	0.322 s
60	0.268 s
80	0.201 s
100	0.161 s

And inverted — the absorption a target demands:

Target RT60	A required at V = 100 m^3
0.30 s	53.67
0.40 s	40.25
0.50 s	32.20
0.60 s	26.83
0.70 s	23.00
0.80 s	20.12
1.00 s	16.10
1.20 s	13.42
1.50 s	10.73
2.00 s	8.05

Read the curve's shape, because it governs every decision that follows: `RT60` is inversely proportional to `A`, so the first few square metres of absorption are worth enormous amounts and the last few are worth almost nothing. Going from A = 8 to A = 16 halves RT60 from 2.01 s to 1.01 s. Going from A = 80 to A = 88 changes it from 0.201 s to 0.183 s. This is the same concavity that governs u in Chapter 7 and it has the same moral: **the second unit is the valuable one, and past a point you are buying glazing bills.**

And the same table across volumes, so you can size any room:

V	A for 0.5 s	A for 0.6 s	A for 0.8 s
30 m^3	9.66	8.05	6.04
50 m^3	16.10	13.42	10.06
75 m^3	24.15	20.13	15.09
100 m^3	32.20	26.83	20.12
150 m^3	48.30	40.25	30.19
200 m^3	64.40	53.67	40.25
300 m^3	96.60	80.50	60.38
500 m^3	161.00	134.17	100.62
1000 m^3	322.00	268.33	201.25
2000 m^3	644.00	536.67	402.50

State A — the as-built hard shell Sealed concrete floor, gypsum board ceiling and walls on studs, 12 m^2 heavy plate glazing. Eight occupants. This is not a caricature; it is what a shell-and-core office looks like before fit-out, and it is what a great many converted spaces look like permanently.

Element	Area	125	250	500	1k	2k	4k
Sealed concrete floor	40.0	0.40	0.40	0.80	0.80	0.80	1.20
Gypsum ceiling	40.0	11.60	4.00	2.00	1.60	2.80	3.60

Element	Area	125	250	500	1k	2k	4k
Gypsum walls	53.0	15.37	5.30	2.65	2.12	3.71	4.77
Glazing	12.0	2.16	0.72	0.48	0.36	0.24	0.24
Occupants x 8	–	1.20	2.00	2.80	3.36	3.68	3.68
Total A		**30.73**	**12.42**	**8.73**	**8.24**	**11.23**	**13.49**
RT60 (s)		**0.52**	**1.30**	**1.84**	**1.95**	**1.43**	**1.19**
`alpha_bar`		0.212	0.086	0.060	0.057	0.077	0.093
Room constant R (m^2)		39.0	13.6	9.3	8.7	12.2	14.9
Critical distance, Q=1 (m)		0.88	0.52	0.43	0.42	0.49	0.54

RT60 of 1.95 s at 1 kHz in a room 8 metres long. The critical distance is 0.42 m — closer than one seat is to the next. Every listener in this room, at every seat, is in the reverberant field. Moving closer to the talker does not help. There is no seat in this room from which you can reliably hear a soft voice.

Note also the spectrum: 0.52 s at 125 Hz against 1.95 s at 1 kHz. The room is *less* reverberant in the bass, because the gypsum-on-studs is doing bass absorption while nothing at all is absorbing mid-frequencies. **This is the signature of an untreated stud-and-board room and it sounds characteristically harsh and clattery.**

State B — the ordinary office fit-out Heavy carpet on the concrete, mineral-fibre suspended ceiling tile over the whole 40 m^2, gypsum walls, glazing unchanged, eight occupants. This is the default fit-out of the last forty years.

Element	Area	125	250	500	1k	2k	4k
Carpet on concrete	40.0	0.80	2.40	5.60	14.80	24.00	26.00
Mineral-fibre ceiling tile	40.0	8.00	16.00	24.00	28.00	30.00	30.00
Gypsum walls	53.0	15.37	5.30	2.65	2.12	3.71	4.77
Glazing	12.0	2.16	0.72	0.48	0.36	0.24	0.24
Occupants x 8	–	1.20	2.00	2.80	3.36	3.68	3.68
Total A		**27.53**	**26.42**	**35.53**	**48.64**	**61.63**	**64.69**
RT60 (s)		**0.58**	**0.61**	**0.45**	**0.33**	**0.26**	**0.25**

Element	Area	125	250	500	1k	2k	4k
Critical distance, Q=1 (m)		0.82	0.80	0.97	1.21	1.46	1.52

State B passes a crude single-number test — 0.45 s at 500 Hz, comfortably inside the ANSI-derived 0.6 s target — and it is nonetheless a badly balanced room. Look at the spectrum:

```
125 Hz : 0.58 s      \
250 Hz : 0.61 s       |  bass ratio (125+250)/(500+1000) = 1.53
500 Hz : 0.45 s       |
 1 kHz : 0.33 s      /
 2 kHz : 0.26 s
 4 kHz : 0.25 s
```

The room is 2.4 times more reverberant at 250 Hz than at 4 kHz. Consonants — the 2-8 kHz band that carries most of the information distinguishing one word from another — are absorbed hard, while the low-mid energy that masks them is left to ring. **A room that over-absorbs treble and under-absorbs bass has, in effect, applied a low-pass filter to every voice in it, and low-pass filtering is exactly what makes speech muddy.**

This is the diagnosis that a single-number RT60 specification cannot produce, and it is why every requirement in §14.14 is stated per band.

State C — the specified deliberation room Wood floor on joists (not carpet), mineral-fibre tile over 20 m^2 of the ceiling with the central 20 m^2 left hard as an early-reflection reflector over the table, 10 m^2 of 100 mm mineral wool with a 50 mm air gap on the end walls, a 9 m^2 filled bookcase as a broadband diffuser-absorber, drapes over half the glazing, eight occupants.

Element	Area	125	250	500	1k	2k	4k
Wood floor on joists	40.0	6.00	4.40	4.00	2.80	2.40	2.80
Mineral-fibre ceiling tile	20.0	4.00	8.00	12.00	14.00	15.00	15.00
Hard ceiling reflector over table	20.0	5.80	2.00	1.00	0.80	1.40	1.80
Mineral wool 100 mm + 50 mm gap	10.0	4.50	8.50	9.50	9.80	9.50	9.00
Bookcase, filled	9.0	2.25	2.70	2.70	2.70	2.70	2.70

Element	Area	125	250	500	1k	2k	4k
Gypsum walls, remaining	28.0	8.12	2.80	1.40	1.12	1.96	2.52
Drape over half the glazing	6.0	0.84	2.10	3.30	4.32	4.20	3.90
Exposed glazing	6.0	1.08	0.36	0.24	0.18	0.12	0.12
Occupants x 8	–	1.20	2.00	2.80	3.36	3.68	3.68
Total A		**33.79**	**32.86**	**36.94**	**39.08**	**40.96**	**41.52**
RT60 occupied (s)		**0.48**	**0.49**	**0.44**	**0.41**	**0.39**	**0.39**
RT60 unoccupied (s)		**0.49**	**0.52**	**0.47**	**0.45**	**0.43**	**0.43**
`alpha_bar`		0.233	0.227	0.255	0.270	0.282	0.286
Critical distance, Q=1 (m)		0.94	0.92	0.99	1.03	1.07	1.08

```
Bass ratio (125+250)/(500+1000) = 1.14        <-- essentially flat
Occupied-to-unoccupied variation = 0.01 to 0.04 s
```

The RT60 varies by 0.10 s across six octave bands. That is a flat room, and flatness is the property that matters — far more than the absolute value.

Note what achieved it, because it is counter-intuitive. **State C has *less* mid-high absorption than State B (36.9 versus 35.5 at 500 Hz is close, but 39.1 versus 48.6 at 1 kHz and 41.0 versus 61.6 at 2 kHz are not).** The design removed the carpet and halved the ceiling tile, and added low-frequency-capable absorption. **The fix for a badly balanced room is often to take absorption away at the top and add it at the bottom, and no single-number specification can express that.**

Note also the hard 20 m^2 ceiling panel over the table. It is not an oversight. It supplies early reflections from the talker to the far listeners, arriving within a few milliseconds of the direct sound, which the ear integrates with the direct sound and which therefore *raise* effective loudness and intelligibility. **Early reflections are the quiet talker's ally; late reverberation is the quiet talker's enemy; and both are made of the same material.** The difference is entirely geometry and timing.

State D — the over-treated room, as a warning Carpet, full mineral-fibre ceiling, 40 m^2 of 100 mm mineral wool on the walls, drapes over all the glazing. Someone was told "add absorption" and did.

	125	250	500	1k	2k	4k
Total A	33.45	59.90	77.65	94.52	104.99	104.65
RT60 Sabine (s)	0.48	0.27	0.21	0.17	0.15	0.15
RT60 Eyring (s)	0.42	0.21	0.15	0.11	0.09	0.09
alpha_bar	0.231	0.413	0.536	0.652	0.724	0.722

```
Bass ratio (125+250)/(500+1000) = 1.99
```

State D is worse than State B, and it cost several times as much. Two failures at once:

1. **It is dead.** At 0.11 s (Eyring) at 1 kHz, the room provides no reverberant support at all. Talkers in dead rooms report having to "push", because none of their own energy comes back; vocal effort rises, fatigue rises, and the quiet talker — who cannot push — is worse off than in a moderately live room.
2. **Its bass ratio has got worse, not better.** All that porous absorption works from 250 Hz up and barely touches 125 Hz, so the *relative* bass excess doubled. The room is dead and boomy at the same time.

And the Sabine/Eyring divergence is 62% at 1 kHz and 78% at 2 kHz. A designer who specified this room using Sabine believed they were building a 0.17 s room and built a 0.11 s room.

14.5 Targets

Anchored on ANSI/ASA S12.60's classroom values, transferred to deliberation with the transfer declared.

Room type	Volume	RT60 target, occupied, 500 Hz-2 kHz	Background, unoccupied
Small meeting / deliberation room	up to 150 m^3	0.40 - 0.60 s	<= 35 dBA (NC-25 to NC-30)
Larger meeting room	150 - 300 m^3	0.50 - 0.70 s	<= 35 dBA
Seminar room / classroom	up to 283 m^3	<= 0.60 s (ANSI limit)	<= 35 dBA (ANSI limit)
Seminar room / classroom	283 - 566 m^3	<= 0.70 s (ANSI limit)	<= 35 dBA (ANSI limit)
Open-plan work area	–	RT60 is the wrong gauge; specify distraction distance and spatial decay	<= 40 dBA
Common room / social space	–	0.6 - 0.9 s (some liveliness is wanted)	<= 45 dBA

And the two requirements that matter more than the absolute number:

```
FLATNESS: max RT60 across 250 Hz .. 4 kHz <= 1.25 x min RT60 across the same bands
 BASS RATIO: (RT60_125 + RT60_250) / (RT60_500 + RT60_1k)  in  [0.9, 1.3]
```

Checking the four states against these:

State	RT60 at 500 Hz	Meets 0.4-0.6 s?	Flatness ratio, 250-4k	Bass ratio	Verdict
A hard	1.84 s	**No**	1.30/0.49… = 5.0	0.48	Fails everything
B office	0.45 s	Yes	0.61/0.25 = 2.4	1.53	**Passes the number, fails the room**
C specified	0.44 s	Yes	0.49/0.39 = 1.26	1.14	Passes (flatness marginal)
D over-treated	0.15 s (Eyring)	**No — too dead**	0.21/0.09 = 2.3	1.99	Fails everything

State B is the entry that justifies this whole section. It passes the single-number test that almost every specification in the world uses, and it is a bad room.

14.6 The two fields, and why the loud voice wins by physics

Now the mechanism. Two things reach a listener: the **direct** sound, which falls 6 dB per doubling of distance, and the **reverberant** field, which is the same everywhere.

```
L_direct(r)  = L_W + 10 log10( Q / (4 pi r^2) )
L_reverb     = L_W + 10 log10( 4 / R )
L_total(r)   = L_W + 10 log10( Q/(4 pi r^2) + 4/R )
```

Take a talker producing 60 dB at 1 m in free field on-axis — nominal "normal vocal effort", the standard reference value — with `Q` = 2, and compare State A with State C at 1 kHz (`R_A` = 8.9 m^2, `R_C` = 53.4 m^2):

r	A: direct	A: reverb	A: total	**A: D/R**	C: direct	C: reverb	C: total	**C: D/R**
0.5 m	66.0	64.5	68.3	**+1.5**	66.0	56.7	66.5	**+9.3**
1.0 m	60.0	64.5	65.8	**-4.5**	60.0	56.7	61.7	**+3.3**
1.5 m	56.5	64.5	65.2	**-8.0**	56.5	56.7	59.6	**-0.3**
2.0 m	54.0	64.5	64.9	**-10.5**	54.0	56.7	58.6	**-2.8**
2.5 m	52.0	64.5	64.8	**-12.5**	52.0	56.7	58.0	**-4.7**
3.0 m	50.5	64.5	64.7	**-14.1**	50.5	56.7	57.6	**-6.3**
4.0 m	48.0	64.5	64.6	**-16.6**	48.0	56.7	57.3	**-8.8**
5.0 m	46.0	64.5	64.6	**-18.5**	46.0	56.7	57.1	**-10.7**

Read the "total" columns first, because they contain the trap. In State A, the total level barely falls with distance at all — 68.3 dB at 0.5 m, 64.6 dB at 5 m. The hard room is LOUDER at the far seats than the treated room is. A sound level meter would report State A as the better room.

Now read the D/R columns. At 3 m, State A delivers -14.1 dB and State C delivers -6.3 dB — a 7.8 dB difference in the ratio that actually determines intelligibility. **Reverberation raises level and destroys clarity. It is loudness without information, and it is the single most common reason a room "sounds fine" on a meter and cannot be worked in.**

Rule: never specify or evaluate a speech room on sound pressure level. Specify RT60 per band and D/R, or measure a clarity index directly.

Why the loud voice wins Now put two talkers in the room. Normal vocal effort is nominally 60 dB at 1 m; raised is 66; loud is 72; shout is 78. A confident, physically larger, or simply more assertive talker at 66 dB and a quiet talker at 54 dB differ by 12 dB — a factor of about 16 in acoustic power, and entirely within the normal range of human vocal output.

Take a listener 3 m away and 40 dBA of HVAC background. Compute the quiet talker's *direct* sound (the only part that carries clarity) against the loud talker's *reverberant* field (which fills the whole room uniformly):

Room	R	Quiet talker's direct at 3 m	Loud talker's reverberant field	Combined noise	**SNR for the quiet voice**
A hard	8.9 m^2	44.5 dB	70.5 dB	70.5 dB	**-26.1 dB**
B office	73.6 m^2	44.5 dB	61.3 dB	61.4 dB	**-16.9 dB**
C specified	53.4 m^2	44.5 dB	62.7 dB	62.8 dB	**-18.3 dB**

And the exclusion radius — how close you must be to the quiet talker before their direct sound exceeds the loud talker's reverberant field at all:

Room	Noise floor from the loud talker	Quiet voice audible within	At +6 dB SNR within	At +15 dB SNR within
A hard	70.5 dB	0.15 m	0.07 m	0.03 m
B office	61.4 dB	0.43 m	0.21 m	0.08 m
C specified	62.8 dB	0.37 m	0.18 m	0.06 m

In State A, while a loud talker is speaking, a quiet talker is intelligible only to someone within 150 millimetres of their mouth. That is inside Hall's intimate zone. In practical terms: they cannot be heard at all.

Two honest notes on this calculation.

Note 1. This is the simultaneous-speech case — the moment of overlap when someone tries to take a turn while another is still speaking. It is not the case where the quiet talker has the floor alone. But **the simultaneous case is exactly the case that determines who gets the floor**, because that is what turn-taking is: a brief overlap in which one party yields.

Note 2. State C's numbers here are slightly *worse* than State B's, because State B has more total absorption at 1 kHz and hence a larger R. The exclusion radius is not the only gauge, and State C wins decisively on flatness, on the tail arithmetic in §14.7, and on early-reflection support. **Report**

the metric that lost as well as the metrics that won. A design that only quotes its favourable numbers is not an engineering document.

14.7 The tail sets the minimum gap — the turn-taking mechanism

This is the most direct link between Sabine's equation and Woolley's finding, and it is pure arithmetic.

Sound decays at `60 / RT60` decibels per second. To take a turn you must be audible over the tail of the previous speaker. So the room imposes a **minimum inter-turn gap**: the time you must wait for the previous speaker's energy to fall far enough for yours to be heard.

Decibels the tail has fallen, after a gap of `t`:

RT60	100 ms	200 ms	300 ms	500 ms	800 ms	1000 ms
0.40 s	-15.0	-30.0	-45.0	-75.0	-120	-150
0.50 s	-12.0	-24.0	-36.0	-60.0	-96.0	-120
0.60 s	-10.0	-20.0	-30.0	-50.0	-80.0	-100
0.80 s	-7.5	-15.0	-22.5	-37.5	-60.0	-75.0
1.00 s	-6.0	-12.0	-18.0	-30.0	-48.0	-60.0
1.50 s	-4.0	-8.0	-12.0	-20.0	-32.0	-40.0
2.00 s	-3.0	-6.0	-9.0	-15.0	-24.0	-30.0
3.00 s	-2.0	-4.0	-6.0	-10.0	-16.0	-20.0

And inverted — the gap required for the tail to fall by a usable margin:

RT60	gap for -10 dB	gap for -15 dB	gap for -20 dB
0.30 s	50 ms	75 ms	100 ms
0.40 s	67 ms	100 ms	133 ms
0.50 s	83 ms	125 ms	167 ms
0.60 s	100 ms	150 ms	200 ms
0.80 s	133 ms	200 ms	267 ms
1.00 s	167 ms	250 ms	333 ms
1.50 s	250 ms	375 ms	500 ms
2.00 s	333 ms	500 ms	667 ms
3.00 s	500 ms	750 ms	1000 ms

Conversation-analytic work reports median turn-transition gaps on the order of 200 milliseconds across a range of languages. **BORROWED AND UNBADGED IN THE CANON — treat the specific figure as indicative, and measure your own room rather than relying on it.** The argument below does not depend on the exact value; it depends only on the gap being short, which is uncontroversial.

Now the mechanism, stated cleanly:

At RT60 = 0.5 s, a 200 ms gap leaves the previous speaker 24 dB down. A quiet talker entering that gap is comfortably above the tail. At RT60 = 2.0 s, the same 200 ms gap leaves

the previous speaker only 6 dB down — and a talker 12 dB quieter than the previous speaker is therefore 6 dB BELOW the tail of the speech they are trying to follow. They are inaudible. They will be talked over. They will experience this as being interrupted, and everyone else will experience it as them not really having said anything.

Nobody in that room did anything wrong. The room did it.

Two secondary consequences follow, and both are observable:

The loud talker is unaffected. At 12 dB above the quiet talker, they clear the tail at 200 ms in a 2 s room too. **Reverberation is not a uniform degradation; it is a filter that selects for vocal power.** The room installs a hierarchy, and the hierarchy is ordered by lung capacity.

The room lengthens turns. If short gaps do not work, speakers learn — without knowing they are learning it — to hold the floor, because releasing it and reclaiming it is expensive. Long turns by few people is precisely the turn-taking-inequality signature Woolley's instrument measures.

From silenced voices to lost ceiling Suppose the room reliably silences three of eight participants. Then the interaction's effective category count is 5, not 8, and Paper D's ceiling `ln N` falls:

`N`	`ln N`	Loss versus `ln 8` = 2.0794
8	2.0794	0.0%
7	1.9459	6.4%
6	1.7918	13.8%
5	1.6094	**22.6%**
4	1.3863	33.3%
3	1.0986	47.2%
2	0.6931	66.7%

Silencing three of eight people costs 22.6% of the room's entire per-interaction surplus ceiling — permanently, in every meeting, for as long as the room stands. And per CANON's honesty clause, `ln N` is a ceiling: the realised loss could be larger, because the three silenced people may have been the ones carrying the non-parallel view. Chapter 11's Parallel Wall and this chapter's Reverberant Room are the same failure arriving from two directions.

A turn-share statistic — and what it is not To measure turn-taking equality, record speaking-time shares `p_1 ... p_n` and compute

```
N_eff_turn = exp( − sum_i p_i ln p_i )
```

BADGE THIS CLEARLY: this is the exponential of the Shannon entropy of the turn distribution, a standard descriptive diversity index. It is NOT the surplus law's `N_eff = 1 + (N−1)u`, and CANON §1.1 is explicit that surplus is not identified with Shannon entropy or mutual information (Paper D). These are two different numbers that happen to be called by similar names. Do not mix them, do not substitute one for the other, and do not report `ln(N_eff_turn)` as a surplus.

With that badge firmly attached, it is a good instrument for the thing it does measure:

Turn-share distribution across 8 people	`N_eff_turn`	as a fraction of 8
Perfectly equal, 12.5% each	8.00	1.000
16/15/14/13/12/11/10/9	7.87	0.983
40/20/15/10/5/5/3/2	5.40	0.675
35/30/15/8/5/4/2/1	5.04	0.630
60/20/10/5/2/1/1/1	3.40	0.425

`N_eff_turn / n` is the field measurement for turn-taking equality, it takes one recording and one line of arithmetic, and it is the number to track before and after acoustic treatment.

14.8 Background noise and the Lombard loop

Background level. Specify it and measure it. NC (Noise Criteria) and NR (Noise Rating) curves specify a maximum level in each octave band rather than a single A-weighted number, which matters because a low-frequency rumble at 40 dBA is far more fatiguing than broadband hiss at 40 dBA.

NC rating	Approximate dBA (nominal, spectrum-dependent)	Typical application
NC-15	~25	Concert hall, recording
NC-20	~30	Very quiet meeting; executive
NC-25	~34	**Deliberation room target**
NC-30	~38	Acceptable meeting room
NC-35	~42	General office; ANSI classroom limit is 35 dBA
NC-40	~47	Busy open plan
NC-45	~52	Too noisy for sustained speech

Signal to noise. For normal-hearing adult native listeners, intelligibility approaches its maximum around +15 dB SNR. **It requires more for children, for non-native listeners, for anyone with hearing loss, and for anyone in an unfamiliar technical vocabulary.** Since a deliberation room will contain some of those people some of the time, +15 dB is a floor and not a comfort.

The vocal effort required to hit +15 dB at 3 m in State C:

Background	Required talker level at 1 m	Vocal effort category
30 dBA	47.4 dB	below normal — easy
35 dBA	52.4 dB	below normal
40 dBA	57.4 dB	approaching normal (60)
45 dBA	62.4 dB	between normal and raised
50 dBA	67.4 dB	above raised (66), toward loud
55 dBA	72.4 dB	**loud (72) — sustained shouting**

At 55 dBA background, a participant must speak at "loud" effort continuously in order to be understood at three metres. Anyone who cannot, or will not, sustain that is out of the conversation. That includes people with respiratory conditions, people recovering from illness, people with soft voices, and people from cultures where raised voices carry meanings they do not wish to convey. **This is an accessibility failure produced by a mechanical services decision.**

The Lombard loop The Lombard effect is the involuntary rise in vocal level in response to background noise. Reported slopes are on the order of 0.3 to 0.6 dB of vocal level per dB of background above roughly 45 dB. **Treat the slope as a reported range, not a constant.**

Now note the structure. In a room with several conversations, each conversation's output is background for the others. That is a feedback loop with loop gain `s`, and Chapter 2's arithmetic applies without modification:

```
total rise = initial disturbance / (1 − s)
```

Lombard slope `s`	Closed-loop gain `1/(1−s)`	A 5 dB seed becomes
0.0	1.00x	5.0 dB
0.2	1.25x	6.2 dB
0.3	1.43x	7.1 dB
0.4	1.67x	8.3 dB
0.5	2.00x	10.0 dB
0.6	2.50x	12.5 dB
0.7	3.33x	16.7 dB
0.8	5.00x	25.0 dB
0.9	10.00x	50.0 dB

At the mid-range reported slope of 0.5, every decibel of avoidable background noise costs two decibels of eventual room level. The system amplifies its own disturbance by a factor of two before settling.

This is **Chapter 1 and Chapter 2's central identity appearing in a completely different domain**: a first-order feedback loop, stable because the gain is below unity, converging to a fixed point by the same contraction argument that makes the storage law converge. The room is a leaky integrator for noise, its DC gain is `1/(1−s)`, and its fixed point is the volume at which everyone in the room is shouting.

Design consequence: the cheapest decibel to remove is the first one, because you remove `1/(1−s)` decibels of eventual level for each decibel of source you eliminate. Fixing a noisy diffuser is worth about twice its nominal value.

14.9 The exclusion arithmetic, and `C_exclusion`

Chapter 12 §12.8 defined `C_exclusion` as the value of contributions the room silenced, and noted that it cannot be measured by asking, because a contribution never attempted leaves no trace. This is the section that computes it from physics.

Procedure — the exclusion calculation.

1. Measure RT60 per band, occupied, and the background level. §14.12 gives the method.
2. Compute `R = A/(1 − alpha_bar)` and `d_c = 0.141 sqrt(Q R)` at 1 kHz.
3. Measure the actual seat separations. Take `r_max` = the greatest talker-to-listener distance.
4. Measure or estimate the room's vocal-level spread. Recording eight participants for one meeting and taking the 10th and 90th percentile of per-talker level at 1 m is adequate; the spread is typically 10-14 dB.
5. Compute, for the quietest decile talker at `r_max`: their direct level, the reverberant field of the loudest decile talker, and the background. Get the SNR.
6. Compute the gap required for the tail to fall 15 dB: `t_15 = RT60/4`.
7. **Count how many participants fall below the audibility criterion.** That count, `k`, is the number of effectively silenced voices.
8. The ceiling loss is `ln N − ln(N − k)`, and the fractional loss is `1 − ln(N−k)/ln N`.

Applied to State A with eight participants, `r_max` = 3.5 m, a 12 dB vocal spread, and the numbers already computed: the quietest three participants have SNR below 0 dB against a loud talker's reverberant field, and `t_15 = 1.95/4 = 488 ms`, roughly 2.4 times a typical turn gap. **k = 3.**

```
ln 8 = 2.0794          ln 5 = 1.6094
ceiling loss = 0.4700 nats = 22.6% of the room's per-interaction ceiling
```

Convert that into the storage law. With `eta` = 0.5 and the room used weekly, the room forgoes `0.5 x 0.47 = 0.235` units of `Q` production per meeting, every meeting, forever. That is `C_exclusion`, and it is now a number rather than a complaint.

Honesty on this procedure, stated plainly. Steps 1-6 are physics. Step 7 — the conversion of an SNR deficit into a count of silenced participants — is a **judgement with a stated threshold**, and different thresholds give different k. State your threshold. Sensitivity-test it. **What the procedure gives you is not a precise loss; it is a defensible order of magnitude computed from measurements rather than from opinions, and that is an enormous improvement over the alternative, which is nothing.**

14.10 Table geometry and sightline

Acoustics and geometry interact, and the acoustic constraint usually bites first.

The governing distance is the longest talker-to-listener path. For a round table of diameter `D`, that is `D`. For a rectangular table of length `L` and width `W`, it is `sqrt(L^2 + W^2)` — corner to corner.

Table	Longest path	In State A (d_c = 0.42 m)	In State C (d_c = 1.03 m)
Round, 1.6 m	1.60 m	3.8 d_c — deep reverberant field	1.55 d_c — marginal
Round, 2.4 m	2.40 m	5.7 d_c	2.33 d_c
Round, 3.2 m	3.20 m	7.6 d_c	3.11 d_c
Rectangular 3.0 x 1.2 m	3.23 m	7.7 d_c	3.14 d_c
Rectangular 4.8 x 1.4 m	5.00 m	11.9 d_c	4.85 d_c

Every practical meeting table puts every listener in the reverberant field. There is no room treatment that makes a 4.8 m table a direct-field environment. What treatment does is raise d_c and therefore reduce how *deeply* into the reverberant field the far seats sit — from 11.9 d_c to 4.85 d_c for the long table, which is a 7.8 dB improvement in D/R at the far corner.

Three geometric rules follow.

Rule 1 — the long rectangular table is an acoustic hierarchy generator. At 4.8 m, corner-to-corner is 5.0 m. The person at one end and the person at the other are 11 dB apart in direct sound relative to a neighbour pair. **The table has a head whether or not anyone sat there deliberately**, because the seats nearest the middle of the long sides are heard by everyone and the seats at the ends are not. Chapter 17 develops the governance consequence.

Rule 2 — reduce the table before reducing the group. A group of twelve at a 2.4 m round table has a 2.4 m longest path. The same twelve at a 4.8 m boardroom table have a 5.0 m longest path and a 6.4 dB worse D/R at the extremes. **The table, not the headcount, was the problem.**

Rule 3 — a hard reflective surface directly above the table is an asset. Early reflections from a hard ceiling panel over the table arrive within a few milliseconds of the direct sound and are integrated with it by the ear. This is precisely what State C's 20 m^2 hard central ceiling panel is for, and it is why blanketing the entire ceiling in absorption (State B, State D) is a mistake. **Absorb the perimeter, reflect the centre.**

Sightline, briefly, because Chapter 17 owns it: visual access to a talker's face contributes materially to intelligibility, and everyone at the table should be able to see everyone else's face without a head turn beyond about 45 degrees. **A round or polygonal table satisfies this automatically; a long rectangle cannot.**

14.11 The cost argument: why this is the cheapest c-factor intervention

The chapter's thesis needs its economics stated, and it can be stated structurally without inventing prices.

Compare three ways of buying turn-taking equality:

Intervention	Cost structure	`delta` (does it decay?)	Requires ongoing attention?	Fails when?
Acoustic treatment	One-time capital; near-zero running cost	Material lifetime, decades	No	Physical damage or refurbishment
Facilitation	Recurring, per meeting, forever	Immediate — every meeting needs it again	Yes, continuously	The facilitator is absent, tired, junior, or outranked
Training in inclusive practice	One-time plus refreshers	Fast; skill and salience decay	Yes	Turnover; and turnover is a `delta` term

Now the storage-law reading. **Facilitation and training are S inputs with high `delta` and a**

recurring C. Acoustic treatment is a C reduction and an N restoration with delta set by concrete and mineral wool. In the language of Chapter 3: treatment moves the intervention from a high-delta store (human attention and skill) into a low-delta store (the building).

That is the entire argument, and it is the same argument Chapter 21 makes about stigmergy: put the state in the medium, not in a coordinator. A facilitator enforcing turn-taking is a coordinator. An absorptive ceiling is a medium. The medium does not get tired, does not get outranked by a senior participant, does not need to be scheduled, and does not leave the organisation.

There is one more asymmetry and it is decisive. **Facilitation works only when someone notices the problem. Acoustics works on the meetings nobody is watching** — the corridor conversation, the unscheduled four people who grabbed the room, the discussion after the chair left. Those are most meetings, and they receive no facilitation at all.

Acoustic treatment of a deliberation room is a one-time capital expenditure that raises turn-taking equality in every meeting held in that room for the life of the building, without administration. Nothing else in this manual has that cost profile.

14.12 The measurement protocol you can actually run

Equipment: a phone with a calibrated (or at least consistently characterised) measurement app, a source capable of a loud broadband burst, and a tape measure. Everything below is achievable in ninety minutes by one person.

M1 — Volume and surface. Measure and compute V and S. Record glazing area separately.

M2 — RT60, unoccupied, per octave band. Use the interrupted-noise method or a balloon burst / starter pistol as an impulse. Take at least three source positions and three microphone positions — nine measurements — and average. Report per band from 125 Hz to 4 kHz. **Report the standard deviation across positions; if it exceeds about 15%, the field is not diffuse and Sabine does not describe the room.**

M3 — RT60, occupied. Repeat with the intended occupancy, or compute it by adding the occupants' sabins to the measured unoccupied A (back-calculate A from the measured RT60, add n x the per-person figure, recompute). **State which you did.**

M4 — Background, unoccupied, services running. Sixty seconds of L_Aeq plus an octave-band spectrum. Take it at the seat, not at the door. **Run this at the time of day the room is used**, not at 7 a.m. when the building is quiet.

M5 — Spatial decay. Place a source at one seat, measure L_Aeq at every other seat with the same stimulus. **A well-treated speech room shows a clear fall with distance; a reverberant room shows almost none.** This is the single most diagnostic five minutes available, and it requires no calibration at all because you are measuring a *difference*.

M6 — Turn-share. Record one real meeting with consent. Compute per-person speaking time and N_eff_turn / n. **Consent is not optional (R14.26).**

M7 — Vocal level spread. From the same recording, per-talker level normalised to 1 m. Report the 10th and 90th percentiles.

M8 — Hand clap. Stand at each seat and clap once. Listen for a discrete echo or a ringing tone. **This costs four seconds per seat and finds flutter echo and focusing that no RT60 number**

will reveal. Do it first.

Report all eight on one page with dates, positions, and equipment. Chapter 24's standard applies.

14.13 The falsification test that Postulate 14.2 requires

Postulate 14.2 chains acoustic treatment to c through Woolley's correlation, and the chain requires a causal direction that is not established. **A programme that acts on Postulate 14.2 owes the following test, and this manual regards omitting it as a failure of the practitioner, not of the theory.**

```
    PREREGISTER, before treatment:

   1. N_eff_turn / n , measured over >= 6 real meetings        [the mediator]
  2. RT60 per band, occupied, and background level          [the intervention]
   3. A group-performance measure agreed in advance, applied
     to >= 6 real tasks the group actually does                [the outcome]

    TREAT.

    MEASURE ALL THREE AGAIN over >= 6 meetings.

    THREE OUTCOMES:

    (a) RT60 improves, N_eff_turn/n improves, performance improves
        -> consistent with Postulate 14.2. Not proof; consistent.

    (b) RT60 improves, N_eff_turn/n improves, performance does NOT
        -> the mediator moved and the outcome did not.
           POSTULATE 14.2 IS NOT SUPPORTED for this group.
           The treatment is still worth having on accessibility and
           fatigue grounds, and should be reported as such and not as
           a c-factor intervention.

    (c) RT60 improves and N_eff_turn/n does NOT
        -> Postulate 14.1 fails for this room. The turn-taking
           inequality here was not acoustic. Look at status, chairing,
           and agenda. Chapters 17 and 22.
```

Outcome (b) is the one nobody reports and it is the one that would tell us the most. Chapter 25 requires the held-out measurement and the shuffled-input control; here the shuffle is straightforward — compare against a matched untreated room over the same period, because the largest confound is that measuring a group changes it.

The Specification

Full requirements for a deliberation room. R14.n. Each names the term of the law it acts on.

Acoustic envelope

R14.1 (acts on `C_effort`, `N`). RT60, **occupied**, in the 500 Hz, 1 kHz and 2 kHz octave bands shall lie between 0.40 s and 0.60 s for rooms up to 150 m^3, and between 0.50 s and 0.70 s for 150-300 m^3.

R14.2 (acts on `C_effort`). Flatness. Across the 250 Hz to 4 kHz octave bands, the maximum RT60 shall not exceed 1.25 times the minimum.

R14.3 (acts on `C_effort`). Bass ratio. `(RT60_125 + RT60_250)/(RT60_500 + RT60_1k)` shall lie between 0.9 and 1.3.

R14.4 (acts on `C_effort`). Lower bound. RT60 shall not fall below 0.30 s in any band from 250 Hz to 4 kHz. A room may be too dead, and State D is what that looks like.

R14.5. Where `alpha_bar` exceeds 0.25 in any band, Eyring's formula shall be used and the calculation shall state which formula was applied in each band.

R14.6. All RT60 figures shall be stated per octave band from 125 Hz to 4 kHz, and shall state whether occupied or unoccupied and at what occupancy. **A single-number RT60 specification shall be rejected.**

R14.7. The 125 Hz figure shall carry a note stating the room's Schroeder frequency `f_s = 2000 sqrt(RT60/V)` and noting that Sabine is unreliable at and below it.

Absorption placement

R14.8 (acts on the diffuse-field assumption). Absorption shall be distributed over at least two non-parallel planes. **A room whose only absorption is a suspended ceiling shall not be certified**, because the horizontal decay path never encounters it.

R14.9. Not less than 30% of the installed absorption, by area, shall be capable of `alpha >= 0.40` at 250 Hz. In practice this means porous absorbers not less than 100 mm thick, or 50 mm with not less than a 50 mm air gap, or a purpose-designed low-frequency device.

R14.10. A hard reflective ceiling zone of not less than 40% of the table's plan area, centred over the table, shall be retained to supply early reflections. **Do not blanket the ceiling.**

R14.11. At least one pair of opposing parallel hard surfaces shall be broken by absorption, diffusion, splay, or furniture, to prevent flutter echo. Verify by hand clap at every seat (M8).

R14.12 (Chapter 13, R13.9). No concave surface with a radius comparable to a room dimension shall face the occupied zone unless demonstrably diffusive, absorptive, or focused away from every seat.

R14.13. Every absorption coefficient used in the calculation shall be traceable to the specific product's tested data, at the intended mounting, tested to a recognised standard. **The nominal table in §14.3 is for design exploration and shall not appear in a construction specification.**

Background noise

R14.14 (acts on `C_effort`, accessibility). Background level, unoccupied with all services running, shall not exceed 35 dBA and shall meet NC-25 to NC-30 in every octave band. Measured at seat height at a seat, at the time of day the room is used.

R14.15. No tonal component shall exceed the adjacent bands by more than 5 dB. Tonal noise is disproportionately fatiguing and drives the Lombard loop harder than its dBA implies.

R14.16. Sound insulation from adjacent spaces shall be sufficient that intrusive speech is inaudible at any seat. **Where the deliberation is confidential, this is a privacy requirement as well as an acoustic one, and it runs in both directions.**

R14.17. Where masking is used, it shall be broadband, steady, and below 42 dBA. Water is a legitimate masking source (Chapter 12 §12.7); an HVAC diffuser is not, because it has a spectrum nobody chose.

Geometry

R14.18 (acts on u). The longest talker-to-listener path shall not exceed 3.0 m for a room intended as a single conversation, and shall be recorded on the drawing.

R14.19 (acts on u). Table geometry shall be round or regular-polygonal. Rectangular tables with an aspect ratio above 2.5:1 shall not be used for deliberation. Chapter 17 gives the governance argument; this chapter gives the acoustic one.

R14.20. Critical distance `d_c = 0.141 sqrt(Q R)` at 1 kHz shall be computed and recorded, and the ratio `r_max / d_c` shall be stated for the furthest seat.

R14.21. Every seat shall have unobstructed sightline to every other seat's face within a 45-degree head turn.

R14.22 (acts on accessibility). At least two seats shall be positioned within 1.5 m of the table centre and shall be reserved-on-request for participants who need the shortest acoustic path. **These shall not be positioned so as to constitute a privileged position; Chapter 17's no-head rule governs.**

Verification and conduct

R14.23. All eight measurements M1-M8 of §14.12 shall be performed at commissioning and recorded with dates, positions, and equipment. **M8 (hand clap) shall be performed first.**

R14.24. RT60 shall be re-measured after any change to furniture, floor covering, glazing treatment or ceiling. A room's acoustics are not a one-time property; removing the drapes changes the specification.

R14.25. `N_eff_turn / n` shall be measured before and after treatment over not fewer than six real meetings each, and reported whether or not it improved.

R14.26 (ethics, CANON §5). No recording of a meeting for turn-share measurement shall be made without the explicit, revocable, cost-free consent of every participant. Recordings shall be processed to speaking-time totals and destroyed. **Any participant may decline, and declining shall not exclude them from the meeting.** A measurement that cannot be declined is not permitted by this manual.

R14.27 (ethics). Acoustic treatment shall not be presented as, or substituted for, a remedy for status-based exclusion. **A treated room removes a physical cause of silence. It does not address a social one, and claiming that it does would let a real problem hide behind a real improvement.**

R14.28 (measurement integrity). The falsification test of §14.13 shall be preregistered before treatment and reported after, including outcome (b) if it occurs.

The Failure Modes

F14.1 — The Reverberant Room

Damages: `N`, `u`, `C_exclusion`. **Description.** RT60 far above target; every seat deep in the reverberant field; the tail sets a minimum turn gap longer than a natural one; turn-taking collapses onto the loudest voices. **Signature.** `N_eff_turn/n` below 0.6. Participants report being "talked over". Meetings run long. A hand clap produces an audible tail. Spatial decay (M5) under 3 dB across the room. **Detects with.** M2, M5, M8, M6. **Fix.** §14.14's specification. Expect `N_eff_turn/n` to move; if it does not, Postulate 14.1 failed here and the cause is social — go to Chapters 17 and 22.

F14.2 — Passed the Number, Failed the Room

Damages: `C_effort`; and the credibility of acoustic specification generally. **Description.** State B. A single-number 500 Hz RT60 inside target, with a bass ratio near 1.5 and 2 kHz over-absorbed. Speech is muddy; nobody can say why; the certificate says it complies. **Signature.** Flatness ratio above 2. Occupants describe the room as "boomy" or "you can hear but you can't understand". **Detects with.** M2 **per band**. A single-number measurement cannot detect this and never will. **Fix.** R14.2 and R14.3. Usually: remove treble absorption, add 100 mm-class low-frequency absorption.

F14.3 — The Dead Room

Damages: `C_effort`, and the quiet talker specifically. **Description.** State D. Over-treatment. No reverberant support, so talkers push; vocal fatigue rises; the quiet talker, who cannot push, is worse off than in a moderately live room. Frequently also bass-heavy, because porous absorbers barely work at 125 Hz. **Signature.** RT60 below 0.30 s at mid frequencies. Occupants report the room is "oppressive", "like a padded cell", "tiring". Bass ratio above 1.6. **Detects with.** M2 per band, with Eyring (R14.5) — **using Sabine here under-reports the deadness by up to 78%. Fix.** Remove absorption from the ceiling centre (R14.10) and from one wall. Add a hard reflector over the table. Counter-intuitive and correct.

F14.4 — Ceiling-Only Treatment

Damages: the diffuse-field assumption, and hence every number you computed. **Description.** All absorption in a suspended ceiling; four hard walls. The horizontal decay path never encounters absorption, so the measured RT60 substantially exceeds the calculated one and the decay curve is two-sloped. **Signature.** Measured RT60 well above calculated. Standard deviation across microphone positions above 15%. Strong flutter between the walls on a hand clap. **Detects with.** M2 with position statistics, and M8. **Fix.** R14.8. Move a third of the absorption onto two non-parallel walls. **The total `A` may not change at all and the room will still improve, because you have fixed the distribution rather than the quantity.**

F14.5 — The Lombard Spiral

Damages: `C_effort`, `C_interrupt`. **Description.** Multiple conversation groups in one acoustic volume, each raising level in response to the others. Loop gain around 0.5 doubles any disturbance. Steady state is a room where everyone shouts. **Signature.** Level rises through the session and does not fall. Occupants leave hoarse. Sound level rises sharply with occupancy, far faster than the `10 log10(n)` you would expect from incoherent addition alone. **Detects with.** `L_Aeq` logged against occupancy over a session. **Fix.** Sub-divide the volume (Chapter 13). Cut the seed disturbance first — **each decibel removed at source removes about two decibels of steady-state level.**

F14.6 — The Services Decision Nobody Reviewed

Damages: `C_effort`, accessibility. **Description.** An HVAC selection, a diffuser type, or a fan speed setting made on a thermal and cost basis alone, raising background to 48-55 dBA and requiring sustained raised-to-loud vocal effort. This is a mechanical engineering decision with a sociological outcome, and the two disciplines usually never meet. **Signature.** Background above 40 dBA at seat height with services running, often with a tonal component. Participants sit forward and lean in. **Detects with.** M4, **at the time of day the room is used**. **Fix.** R14.14 to R14.17. Slower, larger diffusers; attenuators; duct lining; and reviewing the services specification against an acoustic requirement rather than only a thermal one.

F14.7 — The Table That Is Too Long

Damages: `u`, `C_exclusion`. **Description.** A boardroom table with a corner-to-corner path of 4-6 m. The end seats are in a different acoustic world from the middle seats. A hierarchy exists whether or not anyone sat at the head. **Signature.** Turn share correlates with distance from the table's midpoint. End-seat occupants disproportionately silent. `r_max / d_c` above 5. **Detects with.** R14.20's ratio, plus M6 turn-share cross-plotted against seat position. **The cross-plot is the decisive instrument and takes ten minutes. Fix.** R14.19. Replace the table, not the people. This is usually the cheapest single acoustic intervention available and it is often the only one needed.

F14.8 — The Uncounted Exclusion (acoustic case)

Damages: `N`, and the honesty of every metric the room produces. **Description.** Chapter 12's F12.8, specifically. The room silences participants; the silenced cannot report a contribution they never attempted; satisfaction surveys come back fine because the people who could speak were content. **Signature.** A subgroup consistently at the bottom of the turn-share distribution across many meetings, with no other explanation. Late objections from people who "should have said something". **Detects with.** §14.9's exclusion calculation — from physics, not from asking. **Fix.** The specification. And R14.27: **do not let the acoustic fix take credit for a social problem it did not solve, and do not let a social explanation excuse an acoustic problem it did not cause.**

Worked Example

The subject

A weekly departmental decision meeting, eight participants, held in the 100 m^3 room of §14.4 in its as-built State A condition. The department has noticed that the same two people do most of the talking, that decisions get re-litigated, and that three members "never contribute". A management consultant has proposed facilitation training.

Step 1 — Measure (M1 to M8)

M1. 8.0 x 5.0 x 2.5 m. V = 100.0 m^3, S = 145 m^2, glazing 12 m^2.

M8 — hand clap, done first. A clear ringing tail at every seat, and a distinct flutter between the two long walls audible at the table's midpoint. **Four seconds of work, and the diagnosis is already substantially made.**

M2 — RT60, unoccupied, nine positions. Standard deviation across positions 9%, so the field is acceptably diffuse and Sabine applies above the Schroeder frequency. `f_s = 2000 sqrt(1.95/100)` = 279 Hz — **note that this puts both the 125 Hz and 250 Hz bands at or below the Schroeder frequency in this untreated room, and both must be read with caution.**

M3 — RT60, occupied (eight adults).

Band	125	250	500	1k	2k	4k
A (m^2 sabins)	30.73	12.42	8.73	8.24	11.23	13.49
RT60 (s)	0.52	1.30	1.84	1.95	1.43	1.19

M4 — background. 47 dBA at seat height with services running at 14:00, with a 250 Hz tonal component 8 dB above adjacent bands. **R14.15 violated before any other consideration.**

M5 — spatial decay. Source at seat 1, measured at seats 2 through 8: total variation 2.1 dB across the whole table. **The room has essentially no spatial decay, which is the direct observational confirmation of the d_c = 0.42 m figure.**

M6 — turn share, six meetings, with consent per R14.26. Mean shares:

```
0.38, 0.30, 0.13, 0.08, 0.05, 0.03, 0.02, 0.01
N_eff_turn = exp(−sum p ln p) = 4.83
N_eff_turn / n = 4.83 / 8 = 0.604
```

M7 — vocal level spread. 10th percentile 53.5 dB at 1 m; 90th percentile 66.5 dB. Spread 13.0 dB.

Step 2 — Diagnose

Against the specification:

Requirement	Target	Measured	Result
R14.1 RT60 500 Hz-2 kHz occupied	0.40-0.60 s	1.43-1.95 s	**FAIL by 3x**
R14.2 Flatness 250 Hz-4 kHz	max <= 1.25 x min	1.95/1.19 = 1.64	**FAIL**
R14.3 Bass ratio	0.9-1.3	(0.52+1.30)/(1.84+1.95) = 0.48	**FAIL, bass-light**
R14.8 Two non-parallel planes	required	zero absorption installed	**FAIL**
R14.11 Flutter	none	flutter present	**FAIL**
R14.14 Background	<= 35 dBA	47 dBA	**FAIL by 12 dB**
R14.15 Tonality	<= 5 dB	8 dB at 250 Hz	**FAIL**
R14.18 Longest path	<= 3.0 m	3.5 m (2.4 m round table, plus offsets)	FAIL, marginal
R14.20 `r_max/d_c`	recorded	3.5 / 0.42 = 8.3	recorded; very poor

Exclusion calculation (§14.9). Quietest-decile talker at 53.5 dB at 1 m, `r_max` = 3.5 m, loudest-decile at 66.5 dB, R = 8.9 m^2, background 47 dBA. The quiet talker's direct sound at 3.5 m is 42.6 dB; the loud talker's reverberant field is 71.0 dB. **SNR = -28.4 dB.** The gap for the tail to fall 15 dB is `1.95/4` = **488 ms**, roughly 2.4 times a natural turn gap.

Threshold declared: a participant is counted as effectively silenced when their direct sound at `r_max` falls below the room's combined noise floor while another participant holds the floor. On the measured level distribution, **three of eight participants fall below it. k = 3.**

```
ln 8 = 2.0794     ln 5 = 1.6094
ceiling loss = 0.4700 nats = 22.6%
```

And this figure was derived independently of anyone's opinion, from a tape measure, a phone, and Sabine's equation. It matches the department's own observation — "three members never contribute" — but it arrived by a route that cannot be argued with socially.

The consultant's facilitation proposal addresses a problem that the physics says is not primarily behavioural. It may help at the margin. It cannot fix a -28.4 dB SNR.

Step 3 — Specify the treatment

Target State C of §14.4, plus services work.

Work item	Effect
Remove sealed concrete; wood floor on battens over the existing slab, 40 m^2	Adds low-frequency absorption; avoids the carpet treble-heavy trap
Mineral-fibre ceiling tile, 20 m^2, perimeter only	Mid-high absorption where the horizontal path finds it
Retain 20 m^2 hard ceiling over the table (R14.10)	Early reflections; supports the quiet talker

Work item	Effect
100 mm mineral wool with 50 mm air gap, 10 m^2, on the two end walls (R14.8, R14.9)	Low-frequency capable; two non-parallel planes; breaks the flutter
Filled bookcase, 9 m^2, on one long wall	Broadband; also a triangulating object (Chapter 12 §12.6)
Heavy drape over 6 m^2 of the 12 m^2 glazing	Mid-high; adjustable, so N in the lighting category rises too
Duct attenuator and a larger, slower diffuser	Background 47 -> 40 dBA; removes the 250 Hz tone
Replace the 2.4 m round table with a 2.0 m round table	r_max 3.5 -> 2.9 m; satisfies R14.18

Predicted acoustic result, computed:

Band	125	250	500	1k	2k	4k
A (m^2 sabins)	33.79	32.86	36.94	39.08	40.96	41.52
alpha_bar	0.233	0.227	0.255	0.270	0.282	0.286
RT60 Sabine (s)	0.48	0.49	0.44	0.41	0.39	0.39
RT60 Eyring (s)	0.42	0.43	0.38	0.35	0.33	0.33
d_c at Q=1 (m)	0.94	0.92	0.99	1.03	1.07	1.08

Per R14.5, alpha_bar exceeds 0.25 at 500 Hz and above, so **Eyring governs in those bands and Sabine is reported alongside for comparison.** Checking compliance on the Eyring figures:

```
R14.1  500 Hz – 2 kHz : 0.38, 0.35, 0.33 s  -> below the 0.40 s lower target
```

The design as specified overshoots slightly into the dead direction once Eyring is applied. Remove 3 m^2 of the perimeter ceiling tile and leave that area hard. Recomputed at 1 kHz: A = 39.08 - 3(0.70 - 0.04) = 37.10, alpha_bar = 0.256, Eyring RT60 = 0.376 s. Still marginal. Remove 5 m^2 instead: A = 39.08 - 5(0.70 - 0.04) = 35.78, alpha_bar = 0.247, Eyring = 0.392 s.

This is the correction that only appears because R14.5 was obeyed. Had the design been signed off on the Sabine figures alone, the room would have been built 15% deader than intended, in the direction of F14.3. The correction costs 5 m^2 less ceiling tile — it is a saving, not an expense.

Final specification: perimeter ceiling tile reduced to 15 m^2, hard ceiling 25 m^2.

Step 4 — Predicted turn-taking effect

With R at 1 kHz = 53.4 m^2, background 40 dBA, r_max = 2.9 m, and the same 13 dB vocal spread:

```
quiet talker direct at 2.9 m  = 44.3 dB
loud talker reverberant field = 63.2 dB
background                    = 40.0 dB
combined noise                = 63.2 dB
SNR for the quiet voice       = −19.0 dB      (was −28.4 dB)
```

Improvement: **9.4 dB**. Still negative in the simultaneous-speech case — **say so** — but the decisive change is in the tail arithmetic:

```
gap for the tail to fall 15 dB:  RT60/4
    before: 1.95 / 4 = 488 ms       (2.4x a natural turn gap)
    after : 0.39 / 4 =  98 ms       (0.5x a natural turn gap)
```

A natural turn gap now clears the previous speaker by more than 15 dB. Before treatment it did not clear it by 10. That is the mechanism by which a quiet participant recovers the ability to take a turn, and it is a factor-of-five change.

Predicted: k falls from 3 to 0 or 1.

Step 5 — Headroom, before and after

`Q` proxy: standing decisions of this committee currently in force, measured at 12. `delta` = 0.08 per week, measured over a summer recess. `eta` = 0.50 before treatment.

`C` before, in `Q`-equivalent units per week: `C_access` + `C_setup` + `C_effort` + `C_maint` = 0.90, of which the acoustic `C_effort` component is 0.55 (the exclusion loss of §14.9 plus vocal fatigue and the re-litigation caused by decisions taken without three members' input).

```
BEFORE (effective N = 5, because three are silenced):
    eta ln N = 0.50 x 1.6094 = 0.8047
    delta Q  = 0.08 x 12     = 0.9600
    C        =                 0.9000
    h        = 0.8047 - 0.9600 - 0.9000 = -1.0553
```

`h < 0`. Unsustainable. The committee's stock of standing decisions is decaying and no rearrangement of the people prevents it.

After treatment, with `N` restored to 8, `C_effort` reduced by 0.55, and `eta` raised to 0.60 (fewer re-litigated decisions per unit of interaction):

```
AFTER ACOUSTICS:
    eta ln N = 0.60 x 2.0794 = 1.2477
    delta Q  = 0.08 x 12     = 0.9600
    C        = 0.90 - 0.55   = 0.5500
    h        = 1.2477 - 0.9600 - 0.5500 = -0.2623
```

Still negative. Acoustic treatment closed 75.1% of the gap and did not close it. Report that exactly.

The remaining 0.262 comes from `delta Q`. Of the twelve standing decisions, four are dormant — superseded, or naming people who have left — and are consuming 0.32 units per week of maintenance budget while returning nothing. Retire them (Chapter 6's fourth lever, Chapter 23's shedding):

```
ACOUSTICS + SHED Q FROM 12 TO 8:
    eta ln N = 0.60 x 2.0794 = 1.2477
    delta Q  = 0.08 x 8      = 0.6400
    C        =                 0.5500
    h        = 1.2477 - 0.6400 - 0.5500 = +0.0577
```

h > 0. Growth available. r_max = 0.058 per week. Small, but positive, and the arithmetic now says the committee can add one increment of standing order rather than watching its stock decay.

Summary of the four levers:

Configuration	eta	N	ln N	eta ln N	delta Q	C	h
Before	0.50	5	1.6094	0.8047	0.960	0.90	**-1.0553**
Acoustics only	0.60	8	2.0794	1.2477	0.960	0.55	**-0.2623**
Acoustics + shed Q to 8	0.60	8	2.0794	1.2477	0.640	0.55	**+0.0577**
Acoustics + shed Q to 6	0.60	8	2.0794	1.2477	0.480	0.55	**+0.2177**

And the steady state (Paper E, Theorem 3.5), `Q_inf = (eta F(u*) − c*)/delta`, for the treated room at several values of u*:

u*	N_eff = 1 + 7u	F = ln N_eff	eta F	Q_inf
0.20	2.40	0.8755	0.5253	**-0.31**
0.40	3.80	1.3350	0.8010	+3.14
0.60	5.20	1.6487	0.9892	+5.49
0.80	6.60	1.8871	1.1322	+7.28

Compare the same committee in the untreated room, with N effectively 5, eta = 0.50, c* = 0.90:

u*	N_eff = 1 + 4u	F	eta F	Q_inf
0.20	1.80	0.5878	0.2939	**-7.58**
0.40	2.60	0.9555	0.4778	**-5.28**
0.60	3.40	1.2238	0.6119	**-3.60**
0.80	4.20	1.4351	0.7175	**-2.28**

Every entry in the untreated table is negative at every value of u. Per Theorem 3.5, a negative equilibrium means the regime is unsustainable or the linear model must be replaced — and the theorem does not hide this failure mode. In the untreated room, no achievable amount of genuine difference among the eight members produces a sustainable stock of standing decisions. **You could replace every member of that committee with a differently-trained expert and it would still decay.** Treatment moves the equilibrium positive for u* >= 0.4, which is achievable.

Step 6 — The falsification test, preregistered (§14.13, R14.28)

Item	Before (measured)	Prediction	Falsified if
RT60, occupied, 500 Hz-2 kHz	1.84 / 1.95 / 1.43 s	0.40-0.60 s all three bands	any band outside range at commissioning
Flatness, 250 Hz-4 kHz	1.64	<= 1.25	> 1.25
Bass ratio	0.48	0.9-1.3	outside range
Background, unoccupied	47 dBA, 250 Hz tone +8 dB	<= 40 dBA, no tone > 5 dB	> 42 dBA or tone remains
`N_eff_turn / n`, 6 meetings	0.604	>= 0.75	< 0.70
Silenced count k	3	0 or 1	`k >= 2`
Group performance measure	preregistered, 6 tasks	improves	**no improvement -> OUTCOME (b): report that Postulate 14.2 is not supported for this group**
Matched untreated control room	same protocol	no change	control also improves -> the measurement is the cause, not the treatment

The control room is the requirement most often omitted and most often decisive. Measuring a committee's turn-taking for six weeks is itself an intervention; without the control, any improvement is uninterpretable.

Step 7 — Measured at twelve weeks

Metric	Before	After	Target	Result
RT60 500 Hz occupied	1.84 s	0.41 s	0.40-0.60	**PASS**
RT60 1 kHz occupied	1.95 s	0.39 s	0.40-0.60	**marginal fail, 0.01 s low**
RT60 2 kHz occupied	1.43 s	0.38 s	0.40-0.60	**fail, 0.02 s low**
Flatness 250 Hz-4 kHz	1.64	1.18	<= 1.25	**PASS**
Bass ratio	0.48	1.09	0.9-1.3	**PASS**
Background	47 dBA, tone	39 dBA, no tone	<= 40 dBA	**PASS**
Spatial decay across table	2.1 dB	5.4 dB	–	improved

Metric	Before	After	Target	Result
`N_eff_turn / n`	0.604	0.812	>= 0.75	**PASS**
Silenced count k	3	1	0 or 1	**PASS**
Control room `N_eff_turn/n`	0.638	0.650	no change	**PASS — control did not move**

Two bands came in marginally below the 0.40 s lower bound. Report it. It is a 0.02 s miss — within the measurement uncertainty of the method — and the correct action is to record it, note that the room is at the dead end of its target range, and check it again after the room has been in use with its normal complement of coats, bags and papers, all of which add absorption. **Do not adjust the target to match the measurement.**

`N_eff_turn / n` rose from 0.604 to 0.812, a 34% improvement in turn-taking equality, **and the matched control room moved by 0.012, which is noise.** The treatment, not the observation, produced the change.

The group-performance measure showed a small improvement that did not reach the preregistered threshold. **Per the preregistration, this is outcome (b) or a near-miss of it, and it is reported as such: the mediator moved decisively, the outcome moved slightly, and this single case does not establish Postulate 14.2's causal claim.** The treatment is retained on the strength of the accessibility gain, the fatigue reduction, the recovered `ln N` ceiling, and the fact that h went from -1.055 to +0.058 — all of which are worth having on their own account.

What this example demonstrates

1. **A hand clap and a tape measure diagnosed the room in under five minutes**, and Sabine's equation priced the diagnosis in half an hour.
2. **The exclusion of three participants was computed from physics, independently of anyone's report**, which made it un-arguable in a way that a survey never is.
3. **The tail arithmetic — 488 ms to 98 ms — is the mechanism.** It is the sentence to take away: *a reverberant room makes a natural turn gap physically insufficient, and that is how it silences people who are behaving perfectly normally.*
4. **Obeying R14.5 (Eyring above `alpha_bar` = 0.25) changed the design before it was built**, and the change was a saving.
5. **The specification was reported as failing two bands by 0.02 s.** A manual that only reported its passes would be useless.
6. **The causal claim was tested and not confirmed, and that was written down.** Postulate 14.2 remains a postulate. **A framework that cannot fail cannot be trusted, and this one names exactly how it fails.**

Chapter 15 takes the edges and thresholds that this chapter's alcoves and sub-volumes imply, and specifies the transitions between them. Chapter 17 takes the table geometry that §14.10 priced acoustically and prices it as governance.

Chapter 15 . Thresholds, Edges, and Transitions

A building is a set of rooms plus the transitions between them, and the transitions are where most of the cost lives. Chapters 12 to 14 assumed the occupant was already in the room. This chapter covers the part nobody draws: getting in, getting out, and the strip of territory along the wall where — by repeated direct observation — people actually choose to be. Three arguments, which are one argument at three scales.

1. **The edge is where life happens.** People occupy perimeters first, perimeter gets scarcer as rooms grow, and so edge must be manufactured, not merely enclosed.
2. **A transition between two unlike zones reflects.** The larger the mismatch, the more arriving traffic bounces off; a graded transition transmits what a hard junction rejects. **The threshold is not the space before the device; the threshold is the device.**
3. **Every crossing is charged to `C_t`.** Context switching is a line item in the dissipation budget, and it multiplies: crossings per period times cost per crossing. Halve either factor and you have bought headroom without touching η, N or u.

The Law

What is proved

Nothing here is proved about architecture; what is proved is the frame it is fitted into. From Paper E §3, the storage law:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t        S_t = F(U_t) = ln(1 + (N−1) U_t)
```

Cost enters **linearly and negatively** — not discounted, damped or filtered. A unit of C removes a unit of Q at the moment it is spent. Compare η, multiplied against a logarithm and suffering every one of the logarithm's cruelties.

From Paper E, Corollary 3.4 — the sustainability ceiling:

```
a nondecreasing step requires    C_t <= eta ln N − delta Q_t
```

If `C_t > eta ln N − delta Q_t`, no arrangement of the people in the building prevents the stored order from decreasing. You may not hire, diversify or motivate your way out of it.

From Paper E §6.4, headroom:

```
h_t := eta ln N − delta Q_t − C_t
```

This is a chapter about C, and therefore about h. Every metre of graded transition, every alcove cut into a corridor wall, every door that closes is an attempt to move one term of it.

What is postulated

POSTULATE 15.1 (Transition Cost). For a boundary between zones A and B the dissipation charged per period is `C_trans = k * m(Z_A, Z_B, d)`, with `k` crossings per period, `m` the per-crossing cost, `Z_A` and `Z_B` the zones' character vectors (§15.3) and `d` the depth of the device; `m` increases with mismatch and decreases with `d`. **A modelling assumption about form, not a**

theorem; the shape of `m` must be measured in your own system. Both factors are directly observable — count crossings, time recovery — so it is falsifiable where applied.

POSTULATE 15.2 (Impedance Analogy). A graded transition reflects by the same qualitative law as a transmission line: reflection scales with the *ratio* of the two characters, not their difference, and intermediate sections reduce it. **Badged as analogy.** The mathematics is proved for waves on transmission lines and *borrowed* here, doing heuristic rather than deductive work; §15.3 states exactly which part transfers.

POSTULATE 15.3 (Edge Preference). Given free choice, occupants select positions at the perimeter and at interior edges before positions in the open field. **Badged as empirical regularity, not theorem.** Gehl (1971) and Whyte (1980) both report it from systematic observation, and it replicates in your own building in an afternoon. Not a law; what people do.

POSTULATE 15.4 (Threshold Depth Sufficiency). A transition is adequate when a person crossing it can complete the state change the next zone requires *before* arriving in it. **Badged.** The state change is enumerable (§15.4); depth follows from the slowest item on the list.

Honesty clauses carried into this chapter

- Subadditivity holds on the parameter u only. **The triangle inequality fails for geometric triples (Paper A §5).** Do not chain threshold mismatches as though they added like distances; §15.3's step arithmetic is transmission-line arithmetic, not surplus arithmetic.
- As deployed, the surplus functional is a **lower bound** on the ensemble's information. Where this chapter says a threshold protects u, it protects a lower bound on u.
- Paper A makes **no empirical claim**. The law tells you what a saved unit of C is worth once saved; it does not tell you a porch saves one.
- Every measured physical quantity below is computed and the computation shown. Where a number would need field research not done here, this chapter says what to measure instead — explicitly, twice, in §15.7.

The Engineering

15.1 The pattern lineage

Christopher Alexander is this manual's direct ancestor. *Notes on the Synthesis of Form* (1964) argued that design resolves **misfit variables**: you cannot state what a good building is, but you can enumerate the ways one fails its context, and a form is good when the misfits are resolved. That is why Chapter 27 is a failure catalogue and not a virtue catalogue. *A Pattern Language* (1977) gave 253 patterns, each a context, a recurring problem and a configuration that resolves it. The transition and edge patterns:

Pattern	Title	What it is really specifying
112	Entrance Transition	A change of light, level, sound or surface *before* the door
127	Intimacy Gradient	Ordered zones public to private; no zone skipped

Pattern	Title	What it is really specifying
160	Building Edge	The wall of a building as inhabitable territory
179	Alcoves	Sub-rooms within rooms; partial enclosure at seated scale
180	Window Place	The most reliably occupied position in any room

"A City is Not a Tree" (1965) is the load-bearing one. A natural city has **semilattice** structure — units overlap, a shop belongs to a street *and* a neighbourhood *and* a trade — while a planned city is drawn as a **tree**, every unit with one parent and overlaps forbidden. Chapter 8 used this: the tree drives u toward zero, because a strict hierarchy guarantees that any two units at the same level receive their context through the same channel. Thresholds are where the semilattice is realised or destroyed. A room with one door into one corridor is a tree node; a room with a door to the corridor, a glazed panel to the courtyard and a shared alcove to its neighbour is a semilattice node. **The number of ways a room participates is the number of thresholds it has.**

Where this manual departs from Alexander: every pattern here is stated with the term of the storage law it moves. Alexander gives a pattern and a rationale in the language of life and wholeness; this manual gives a pattern, the parameter it changes, and the arithmetic that prices it — a narrower and more brittle claim, deliberately, because a claim you can check is worth more to an engineer than one you can only agree with.

Pattern family	Parameter	Mechanism
Entrance transition, airlock, bent entry	`C` down	Reduces per-crossing cost `m`
Door that closes, screened alcove	`C` down	Reduces crossing count `k`
Building edge, window place, pocket	`eta` up	Occupied time where encounters become durable
Intimacy gradient	`C` down, `u` kept	Prevents the wrong-mode encounter
Semilattice adjacency	`N` up	More genuine categories of participation
Desire-line paving	`eta` up	Converts an observation into durable structure

15.2 The edge effect, and the arithmetic of edge scarcity

Gehl's *Life Between Buildings* (1971) divides outdoor activity into three kinds:

- **Necessary** — going to work, waiting for the bus, carrying the shopping. These happen in nearly any condition; quality of space barely affects their frequency.
- **Optional** — sitting down, standing to look, walking with no destination. Only when invited.
- **Social** — everything requiring others, including the passive contact of seeing and hearing people you do not speak to.

Optional activity is the observable proxy for η. Necessary activity happens whether or not the space is any good and tells you nothing; optional activity is discretionary presence, the substrate from which durable relation and decision get made.

The edge finding is the second half. Under systematic observation people take up positions at edges — walls, arcade columns, steps, planters, the boundary between paving and grass — before positions in open ground. The mechanism usually offered is prospect and refuge: back covered, view forward and wide, exit one step away, and *partial* engagement, which is the only kind most people want most of the time. **That mechanism is a hypothesis; the occupancy pattern is an observation, and you can confirm it in your own atrium before lunch.**

The engineering consequence is arithmetic. Edge is perimeter; perimeter scales with the first power of linear dimension and area with the second, so the good stuff gets scarcer as the room gets bigger:

```
square room, side s:  edge density  P/A = 4s / s^2 = 4/s  metres of edge per square metre
```

Side (m)	4	6	8	10	15	20	30
Area (m²)	16	36	64	100	225	400	900
Edge density (m/m²)	1.000	0.667	0.500	0.400	0.267	0.200	0.133

A 30 m square hall has one-seventh the edge per square metre of a 4 m room. **This is the structural reason large open-plan floors feel worse than their area predicts: the amenity people select for has been diluted sevenfold while the rent per head stayed constant** — and furniture quality does not compensate, because the deficit is topological.

The fix is not small rooms but to **manufacture interior edge**, and the arithmetic is unusually favourable: an interior island — bookcase run, planter bank, service core, low wall, level change, storage bank — contributes its full perimeter while consuming only its footprint. For a 10 x 10 m room (A = 100 m², P = 40 m):

Intervention	Edge added	Gain	Floor consumed	Edge per m² consumed
1 island 4.0 x 4.0 m	+16.0 m	+40%	16.0 m² (16.0%)	1.00
2 islands 2.0 x 4.0 m	+24.0 m	+60%	16.0 m² (16.0%)	1.50
4 islands 1.2 x 3.0 m	+33.6 m	+84%	14.4 m² (14.4%)	2.33
6 islands 0.9 x 2.4 m	+39.6 m	+99%	13.0 m² (13.0%)	3.06

Six thin islands nearly **double** the room's edge for 13% of its floor; one fat island buys 40% for 16%.

Edge yield is maximised by many thin obstructions, not few thick ones. The limit is circulation clearance: below about 0.9 m clear between islands the gain is cancelled by congestion, itself a C term.

On a 12 x 18 m room (A = 216 m², P = 60 m) the same six islands add 66% rather than 99%, because the base perimeter is larger relative to the intervention. **Edge manufacture is most valuable exactly where edge is scarcest, which is the deep plan** — and that is where it is most often omitted, because islands read as lost area on the same spreadsheet that produced the deep plan.

15.3 Thresholds as impedance matching

Now the analogy, badged as Postulate 15.2, made precise about what transfers.

On a transmission line, a wave in a medium of characteristic impedance Z_1 meeting a medium of impedance Z_2 is partially reflected:

`Gamma = (Z_2 − Z_1) / (Z_2 + Z_1)` `power reflected = |Gamma|^2, transmitted = 1 − |Gamma|^2`

Ratio Z_2/Z_1	2	4	6	10	100	1000
Γ	0.333	0.600	0.714	0.818	0.980	0.998
Power transmitted	**88.9%**	**64.0%**	**49.0%**	**33.1%**	**3.9%**	**0.4%**

Two structural facts, and these are the ones that transfer:

Fact 1 — reflection depends on the RATIO, not the difference. A step from 1 to 2 and one from 500 to 1000 reflect identically; absolute character is irrelevant, only the jump.

Fact 2 — intermediate sections reduce the worst junction. Take the brutal 1 -> 100 case and insert one section at the geometric mean `Z_t = sqrt(Z_1 Z_2) = 10`. The junctions 1 -> 10 and 10 -> 100 now have *identical* coefficients of 0.818 — the geometric mean is the value that equalises them. Subdivide further with a constant ratio per step `r = (Z_2/Z_1)^(1/n)`:

Steps n	1	2	3	4	5	6
Ratio per step	100.00	10.00	4.64	3.16	2.51	2.15
Γ per junction	0.980	0.818	0.646	0.520	0.431	0.366
Reflected per junction	96.1%	66.9%	41.7%	27.0%	18.5%	13.4%

Worst-junction reflection falls from 96% to 13% by dividing the same total change into six graded steps; in the limit the tapered line reflects almost nothing. **A threshold is a matching section. That is the whole of the idea, and the whole of what this analogy is for.**

What does NOT transfer, stated flatly. The transmission-line result depends on wave phase: the quarter-wave transformer achieves *zero* reflection only where the section is a quarter wavelength long, and tapered designs are broadband because they trade that null for tolerable performance across a band. Nothing in a corridor has a phase. **Do not compute a corridor length from a wavelength; there is no wavelength.** Two sentences transfer — mismatch is a ratio, subdivision reduces it — and anyone extending the analogy past them has left engineering for decoration.

The zone character vector To use the analogy you need `Z` — a vector of ordinal attributes, each rated 1 to 5, each independently observable.

Attribute	1	3	5
Sound level	library quiet	conversational	machinery, crowd
Pace	seated, hours	task, minutes	transit, seconds
Formality	domestic	working	ceremonial, on-record
Illuminance	low, warm, local	mixed	high, uniform, cool
Ownership	one named person	a known group	anyone
Attention	deep, single-thread	interruptible	scanning
Body state	shoes off, seated	standing, coat on	moving, loaded
Observability	unseen	seen by peers	seen by strangers

Rate both zones. The **mismatch index** is `M = max_i |Z_A[i] − Z_B[i]|`, with `M_bar` the mean of the same differences.

`M` is what matters: a transition can be graded on seven attributes and violent on the eighth, and the eighth is what people feel. Treat `M >= 3` as requiring a device whatever `M_bar` says.

Design rule 15.A — one step per threshold. No threshold may change any attribute by more than 1 point on the 5-point scale; if `M = 4` you need at least three intermediate zones. This is the direct translation of the taper table and the most useful single rule in the chapter.

Worked instance — street to reading room:

Attribute	Sound	Pace	Formality	Illum.	Ownership	Attention	Body	Observ.
Street	5	5	3	5	5	5	5	5
Reading room	1	1	2	2	2	1	1	2
Gap	4	4	1	3	3	4	4	3

`M = 4`. A single door here is the 1 -> 100 junction, and the bouncing is observable — hesitation at the door, entries followed by exits within thirty seconds, standing in the entry looking for permission, street volume for two sentences then an over-corrected silence, or simply not going in. Every one is a reflection chargeable to `C`.

The graded version inserts four intermediate zones, one step each:

```
street        forecourt      porch        lobby       anteroom     reading room
sound   5  -->     4     -->     3     -->     2     -->     2     -->      1
```

```
pace   5  -->   4   -->   3   -->   2   -->   2   -->   1
attn   5  -->   4   -->   3   -->   3   -->   2   -->   1
body   5  -->   4   -->   3   -->   2   -->   2   -->   1
own    5  -->   5   -->   4   -->   3   -->   2   -->   2
```

Each junction now has `M = 1`, and worst reflection collapses from near-total to marginal without a single change to the reading room itself. Notice what the intermediate zones physically are: a forecourt changes surface and width; a porch, overhead and light; a lobby, acoustic and ownership; an anteroom, pace and observability. **Each changes one variable, and that is what makes it an intermediate zone rather than a small room.**

15.4 The entrance sequence: what a threshold has to do

A threshold is a machine with a duty list, and its depth is set by the slowest item on it:

1. **Decelerate**, from transit speed to standing. 2. **Reorient** — find the next direction, from signage or geometry. 3. **Shed**: coat, umbrella, shoes, bag, load, phone call. 4. **Be seen** by whoever is inside, before you are among them. 5. **See** — assess the room before committing to it, the reciprocal of item 4 and the one most often omitted. 6. **Choose**: go in, wait or leave, none of the three conspicuous. 7. **Be greeted**, or deliberately not, but by architecture rather than by accident. 8. **Store** the shed items where they will still be at the end.

Every classical entrance device solves some subset of this list, and they are worth naming because they are convergent evolution and therefore evidence:

Device	Duty items served
Japanese **genkan** — floor-level change, shoes removed and stored	1, 3, 8; marker of 4 and 6
Japanese **engawa** — covered verandah, neither in nor out	1, 5, 6; partial commitment
Roman **vestibulum -> fauces -> atrium** — forecourt, throat, hall	a literal three-section taper
Christian **narthex** — antechamber before the nave	4, 5, 6; gradient of belonging
Bent entrance (Persian, Arab, Ottoman) — never on axis	1, 5, 8 at once
English **porch with bench** (Alexander Pattern 242)	1, 6, 7
Industrial **airlock** — two doors never open together	1–8, made hardware

None of these is decoration. Each is a matching section.

Deceleration sets the minimum depth Comfortable walking speed on the flat is about 1.2 to 1.4 m/s for adults, slower for children, the elderly and the loaded. A comfortable voluntary deceleration — one that does not read as a stumble or a halt — is on the order of 0.5 m/s². Then

```
d = v^2 / (2a)          t = v / a
```

v (m/s)	1.2	1.4	1.4	1.4	1.6
a (m/s²)	0.5	0.3	0.5	1.0	0.5
Stopping depth d (m)	1.44	3.27	**1.96**	0.98	2.56
Time t (s)	2.40	4.67	2.80	1.40	3.20

A threshold that expects people to stop must be at least 2.0 m deep in the direction of travel. Below that the stop is achieved by the walker's own braking effort rather than by the space, and registers as being brought up short. This is why a 1.2 m vestibule feels wrong even when correctly sized for the door swing: the door was designed and the deceleration was not.

Dwell time sets the useful depth Deceleration is the floor. The threshold's real job is to hold a person long enough for the state change, and the currency of that is time, not metres — transit time `t = d / v`:

Depth d (m)	1.5	2.0	2.5	4.0	6.0
at 1.4 m/s	1.07 s	1.43 s	1.79 s	2.86 s	4.29 s
at 0.7 m/s	2.14 s	2.86 s	3.57 s	5.71 s	8.57 s
at 0.4 m/s	3.75 s	5.00 s	6.25 s	10.00 s	15.00 s

At full walking pace a 2 m threshold buys 1.4 seconds — one breath, enough for item 1 and nothing else. To buy the four or five seconds items 2, 4 and 5 need there are two options: 6 m of depth, or **slow the walker down**, at which point 2.5 m buys 6.25 s.

This is the case for the **bent entrance**, and the reason it recurs across unrelated traditions. A 90deg turn cannot be taken at 1.4 m/s; it forces a drop to roughly 0.7 m/s around the corner, converting 1.5 m of path into about 2.1 s instead of 1.1 s, while delivering item 5 (the interior revealed progressively, as you turn) and item 4 (seen approaching, before you are present). **The turn is worth more than the metres it costs, which is why every deep transition in the traditions bends.**

The other dimensions A threshold a person can pause in must let another pass, or the pause is an obstruction and the obstruction is a cost.

```
stopped person, footprint depth  0.60 m     clear width, one person passing  0.75 m
body depth (chest to back)      0.45 m     clear width, two abreast       1.20 m
shoulder breadth        0.45 – 0.50 m     door leaf 0.9 m -> swing radius  0.90 m

holds one paused person AND permits single-file passage:  0.60 + 0.75 = 1.35 m clear width
inward-swinging 0.9 m leaf, holding one paused person:   0.90 + 0.60 = 1.50 m landing depth
holds a conversing PAIR at Hall personal distance and
  still permits passage:         0.45 + 1.20 + 0.45 = 2.10 m, + 0.75 m to pass = 2.85 m
```

The last figure matters most for workplaces, clinics and schools, because the conversation at the entrance is a large fraction of the optional activity in the whole building. **A 2.85 m entrance lobby is not generous. It is the minimum width at which two people can talk without becoming an**

obstacle — and if they become an obstacle they stop talking, and the optional activity that was your η proxy disappears.

15.5 The intimacy gradient

Alexander's Pattern 127 states the sequence; Hall's *The Hidden Dimension* (1966) supplies the metric — proxemic bands for adults in the culture Hall studied:

```
intimate    0.00 – 0.45 m
personal    0.45 – 1.20 m
social      1.20 – 3.60 m
public          > 3.60 m
```

These are not universal constants and Hall did not claim they were; they are culturally and individually variable. **What is reliable is the ORDERING and the existence of the bands, not the exact boundaries.** Verify them in your own population, and expect the personal/social boundary in particular to move.

The building-scale gradient runs on the same logic:

Zone	Who may enter	Without asking?	Pace	Typical device
Public	anyone	yes	transit	street, forecourt
Semi-public	anyone with business	yes	slowed	lobby, café, front desk
Common	members	yes	task	corridor, kitchen, shared table
Semi-private	members with reason	knock or hover	task	team room, open alcove
Private	invited	ask	deep	office, study, closed room
Intimate	chosen	ask	any	one's own desk, one's own chair

Design rule 15.B — no zone may be skipped. A door opening directly from public to private is the 1 -> 100 junction and both sides pay. The receptionist whose desk faces the street door absorbs a four-step mismatch all day, and the cost appears in the turnover rate, a δ term.

Design rule 15.C — the gradient must be legible without signage. If a visitor needs a sign to know they have left the common zone, the gradient is not doing its work; the sign is. Signs are administered — made, mounted, updated, enforced — and geometry is not. Change the floor surface, ceiling height, light level, width or direction; each is read instantly by everyone, including people who cannot read the sign.

15.6 Desire lines as measured gradient

Book One, Node 02 taught the operator to read a desire line. Here it becomes an instrument.

A desire line is the physical trace of accumulated route choice: the worn diagonal across the lawn, the polished floor patch, the chair that has migrated three metres. It is a *measurement*

with properties no survey has. It is **unprompted** — nobody was asked, so nobody performed an answer, which is the entire problem with asking people about their building. It is **weighted by frequency**: the depth of the trace integrates over crossings, making it literally an accumulator, in the same family as `Q`. It **records cost, not preference** — people route around cost, and a desire line marks where the designed path was dearer than the undesigned one. And it is free and continuously updated.

Protocol 15.1 — the desire-line survey.

1. Photograph every floor and ground surface from a taped, fixed camera position, same time of day, weekly for six weeks, so the images register. Record wear differentially: floor polish, carpet nap direction, grass, dust, and snow or dew — the best desire-line media in existence, free for a few days a year.
2. Photograph every room at end of day *before* it is tidied. Chair positions are a desire-line medium, as are mugs, cables, coats and laptops left in situ.
3. Mark handle wear and scuffs at 0.9 m (bag corners) and 0.15 m (feet). Log door-propping: a habitually propped door is a threshold whose k is high enough that the occupants chose to pay in noise rather than in operations. That is a datum.
4. Overlay the six images. Traces in all six are structural; traces in one or two are events.

Protocol 15.2 — what to do with it. Pave the structural desire lines. Paving converts an observation into infrastructure — Chapter 4's η-raising move — and stops the maintenance cost of repaired grass and re-polished floor, a direct C reduction.

Honest caution. Paving destroys the instrument that produced the trace. Pave every one and a continuously-updating measurement becomes a fixed plan, with nowhere for the next generation of route preferences to write itself — Goodhart's Law arriving through the floor. Retain one unpaved surface beside every paved route. **A building with no soft surfaces left has lost the ability to tell you it is wrong.**

15.7 Context switching as a C term

The chapter's most consequential arithmetic, and where two figures must be measured rather than asserted. The per-period switching cost:

```
C_switch = k * m
```

- `k` = number of involuntary context switches per period per person
- `m` = the **resumption lag**: elapsed time from the end of the interruption until the person has regained the working state they were in before it

Both are observable — `k` a count, `m` a timed measurement. On an eight-hour day (480 min):

k / day	m = 2 min	m = 5 min	m = 10 min
5	10 min (2.1%)	25 min (5.2%)	50 min (10.4%)
10	20 min (4.2%)	50 min (10.4%)	100 min (20.8%)
20	40 min (8.3%)	100 min (20.8%)	**200 min (41.7%)**
40	80 min (16.7%)	200 min (41.7%)	400 min (83.3%)

Twenty interruptions a day at a ten-minute resumption lag consumes 42% of the working

day before any work has been done wrong. The product is bilinear: halving either factor halves the cost, and the two levers are different engineering problems. k is a threshold problem — filters, doors, distance, visibility control. m is a task-state problem — externalised working state, checkpoints, a written next-action, an artifact that holds the context so the person need not.

MEASUREMENT REQUIRED, NOT ASSERTED. Field studies of interrupted knowledge work report substantial resumption lags, and the finding that recovery costs materially more than the interruption itself is well established — Gloria Mark's field research is the standard reference. But magnitudes vary enormously with task and interruption type, with whether the interruption is self-initiated, and with how "resumed" is operationalised. **This manual supplies no number for m: any figure quoted as universal has been detached from its measurement definition.** Measure it:

Protocol 15.3 — measuring m. For a task with a natural progress metric (lines written, cases cleared, tests passed), record it in five-minute buckets for at least ten person-days, logging every interruption's start and end. Take the mean rate in the twenty minutes before each interruption; find the elapsed time after it ends at which the rate returns to within 90% of that baseline. That is m. Report median and interquartile range — the distribution has a long right tail.

Protocol 15.4 — measuring k. A tally card taped to the desk, one mark per involuntary switch, for two weeks. Self-report undercounts, but *consistently*, which keeps before-and-after comparison valid even when the absolute level is not. Do not use telemetry alone: it misses the person leaning over the partition, usually the dominant channel.

Distance as a filter The threshold's role in k is to be a filter with tunable admittance, roughly in order of strength:

Device	Effect on k	Cost	Reversible?
Physical distance	strong	high	no
A door that closes	strong	low	instantly
Screen at 1.35–1.45 m	moderate	low	yes
Screen at 1.20 m	weak	low	yes
Orientation to traffic	moderate	zero	yes
Signal convention	moderate	zero	needs upkeep
Timed availability windows	strong	zero	yes

Two of these deserve comment.

Distance. Thomas Allen's studies of R&D organisations established the regularity now called the Allen curve: the probability of regular technical communication falls steeply with physical separation and is near its floor within a few tens of metres. The shape is the point, not the coefficients — a steep decay, so the first ten metres do most of the work. **This cuts both ways: the same distance that suppresses the interruptions you did not want suppresses the encounters that were your η.** A blunt filter, reducing k and eta together — first for strength, last for advisability.

Screen height. Seated eye height is roughly 1.15–1.25 m, standing 1.55–1.70 m. A screen at 1.20 m sits at seated eye level: it blocks a seated forward view but is trivially seen over by anyone standing, delivering seated privacy plus standing awareness. A screen at 1.35–1.45 m blocks the

seated view completely and makes a standing person deliberately look over, converting a glance into an act. **The 15 cm between those heights is the difference between a room where you are interruptible by accident and a room where interrupting you is a decision** — the cheapest single k intervention in this chapter.

15.8 What a threshold costs, and whether it pays

A five-zone entrance sequence for a medium building — forecourt 3.0 x 4.0, porch 2.0 x 3.0, vestibule 2.4 x 2.4, lobby 4.0 x 5.0, anteroom 2.4 x 3.0 — totals **51.0 m²**: 8.5% of a 600 m² floor, 5.7% of 900 m², 4.2% of 1200 m². Four to nine per cent of net area, and the first thing value-engineered out of a project, because on a plan it reads as circulation and circulation reads as waste.

Against that, the recovered capacity. For 30 people, suppose threshold work moves k from 22 to 6 per day at a measured `m` of 7 min:

```
before 22∗7 = 154 min/day ; after 6∗7 = 42 min/day ; saved 112 min/day = 1.87 h/person/day
30 people ∗ 220 working days ∗ 1.87 h = 12,320 person-hours/yr = about 7.7 FTE-years per year
```

Fifty-one square metres of transition against twelve thousand person-hours a year. Even at a third of the assumed `m` it is not close. The honest qualifications: `k = 22 -> 6` is not achieved by architecture alone, and the protocol half must be maintained, a δ problem rather than a build; not all recovered time converts to output, because the conversion is η and η is not 1; and area is capital while time is operating, so state both rather than netting them silently.

15.9 Where a threshold must not be

Prohibition 1 — a threshold decelerates; it does not deny. Every device here must be crossable by anyone who chooses. The bent entrance slows you; it does not stop you. The closed door has a handle on both sides. The screen can be looked over. **A threshold that removes a person's ability to enter is not a threshold, it is a gate, and gates are outside the scope of this manual** (CANON §5.4, and its symmetric case). The distinction is operational: a threshold raises the cost of crossing and leaves the choice with the crosser; a gate moves the choice to a third party.

Prohibition 2 — no threshold on an egress route. Deceleration devices, bends, level changes and closable doors are obstructions under evacuation conditions; each must be checked against the applicable life-safety code, and where a conflict exists code wins without discussion. This manual supplies no egress widths, travel distances or door hardware: those are jurisdictional and they change. **Get the code from the authority having jurisdiction.**

A third, softer caution: filters reduce k *selectively*, and the selection is not neutral. A door that closes filters hardest against those with the least standing to knock — the new, the junior, the visiting — so silencing exactly the members whose views were least correlated with the rest raises η by cutting u, which is Chapter 11's failure mode arriving through the door frame. The counter is architectural: **pair every closable private zone with an open, low-commitment, edge-type position anyone may occupy without asking.**

The Specification

Where a figure is jurisdictional or population-specific it says so, and says what to measure.

S15.1 — Mismatch survey. Rate both adjoining zones on the eight-attribute character vector of §15.3; record `M` and `M_bar`. A design not beginning with these two numbers is not a design.

S15.2 — Step limit. No threshold shall change any character attribute by more than 1 point on the 5-point scale; where `M = j`, provide at least `j − 1` intermediate zones.

S15.3 — Deceleration depth and dwell. Any threshold at which occupants are expected to stop shall provide at least **2.0 m** of depth in the direction of travel ($d = v^2/2a$ at v = 1.4 m/s, a = 0.5 m/s²); where 2.0 m is unavailable, substitute a direction change of at least 90deg. Compute $t = d/v$ against the duty list of §15.4; where items 2, 4 and 5 are required, target $t \geq 4$ s, achieved by lowering v rather than by adding depth.

S15.4 — Landing depth and pause-and-pass width. On the swing side of an inward-opening 0.9 m leaf, landing depth not less than **1.50 m** (0.90 swing radius + 0.60 standing footprint). Clear width not less than **1.35 m** where one person may pause, and **2.85 m** at any entrance where two are expected to converse (0.45 + 1.20 + 0.45 for the pair at Hall personal distance, plus 0.75 m passing clearance). Latch-side clearances for accessible approach are jurisdictional; obtain them from the applicable standard.

S15.5 — Legibility without signage. Every zone transition shall be marked by at least two of: floor surface change, ceiling height change, illuminance change of 2x or more, width change of 30% or more, direction change of 45deg or more. Signage may supplement, never substitute.

S15.6 — Edge density. Any occupiable room over 60 m² shall achieve total perimeter (including interior islands) over floor area of not less than **0.40 m/m²**, preferring many thin islands: at equal floor consumption six islands of 0.9 x 2.4 m yield 2.5 times the edge of one 4 x 4 m.

S15.7 — Island clearance and occupiable edge. Not less than **0.90 m** clear between interior islands, **1.20 m** on any route two people use at once. Not less than **60%** of the edge counted under S15.6 shall be occupiable — beside a sittable, leanable or standable position with at least 0.60 m clear depth. Edge that cannot be inhabited does not count.

S15.8 — Window place and alcove sizing. Every room with a window shall have at least one seated position with back or side to a wall, face to the window, within **1.5 m** of the glazing (Alexander's Pattern 180, as a requirement). Two-person alcove: clear depth **2.10 m** minimum (0.45 + 1.20 + 0.45), clear width **1.50 m**. Four-person alcove seated on a 1.8 m circle: clear diameter **2.70 m**.

S15.9 — Screen heights. For seated privacy with standing awareness, **1.20 m**. To block the seated view so interruption becomes a deliberate act, **1.35–1.45 m**. Do not specify heights between 1.25 and 1.30 m; they deliver neither condition.

S15.10 — Closable boundary. Every zone rated 1 or 2 on attention mode shall have a physically closable boundary, operable by its occupant from inside, without tools and without permission.

S15.11 — Paired invitation. Every closable private zone under S15.10 shall be paired with an open, unassigned, edge-type position within 15 m that anyone may occupy without asking. Filters and invitations are specified together, or not at all.

S15.12 — Desire-line survey. Run Protocol 15.1 for six weeks before any layout change and six months after, retaining one unpaved or unfixed surface beside every paved route.

S15.13 — Switching measurement. Measure k (Protocol 15.4) and m (Protocol 15.3) before and after any intervention; report C_switch = k.m as a fraction of the working period. **No threshold intervention may be claimed successful on satisfaction scores alone.**

S15.14 — Egress precedence and crossability. Check every threshold device against the applicable life-safety and accessibility codes; where this specification conflicts, code governs, and no deceleration device, bend, level change or closable boundary shall sit on a required means of egress except as the code expressly permits. Every threshold shall be crossable in both directions by any person who chooses, without a third party's permission; devices transferring the crossing decision to a third party are gates, outside this specification.

S15.15 — Entrance area budget. Allow **4% to 9%** of served net floor area for the entrance sequence, sized by mismatch, as a named line item at the first cost plan — so later removal is a visible decision rather than a silent one.

The Failure Modes

THE CLIFF EDGE. A single door between zones with M >= 3 — classically the street door opening straight into a working room. *Damages* C (high m) and eta (most approaches abort). *Signature:* hesitation over two seconds at the door; entries followed by exits within thirty seconds; people standing just inside, scanning; highest turnover on the floor among the staff nearest the door. *Fix:* intermediate zones per S15.2, or — where no area exists — a 90deg bend and a surface change, which costs nothing and delivers deceleration plus reveal.

THE CORRIDOR WITH NO POCKET. A route of uniform width with no recesses, so any pause is an obstruction. *Damages* eta (kills the incidental encounter, the cheapest η source in a building) and C (congestion). *Signature:* corridor conversations under fifteen seconds, ending with one party stepping aside; meetings held standing in doorways; the corridor empty except between scheduled events. *Fix:* every 12–15 m, widen by 1.2 m over a 2.4 m length with something to lean on — about 3 m² per pocket, the highest-yield η intervention per square metre anywhere.

THE DEEP PLAN. Edge density below 0.2 m/m². *Damages* eta: the positions people select for do not exist in sufficient number. *Signature:* perimeter positions occupied ~100% of hours, interior under 40%, regardless of assignment; a doughnut in the occupancy data. *Fix:* manufacture interior edge per S15.6. If S15.6 stays unreachable the room is the wrong shape and no furniture saves it — subdivide.

THE REFLECTED VISITOR. A threshold physically deep enough but delivering none of duty items 4, 5 or 7. *Damages* C, and N indirectly: the people willing to cross narrow to those already sure of welcome. *Signature:* visitors phone from the lobby; "reception" appears in wayfinding but no person is visible from the door. *Fix:* a sightline from the pause point into the destination and a reciprocal one out; one occupied sitting position visible from the entry.

THE PAVED MEADOW. Every desire line paved; no soft surfaces left, so no way to record a route preference — a Goodhart failure in which the measure became the plan and stopped measuring. *Signature:* no wear patterns, nothing moved in two years, every route correct and none informative. *Fix:* restore soft sensors per S15.12.

THE FORTIFIED THRESHOLD. A threshold that has become a gate: crossing needs a third party's assent. *Damages* u (filters hardest against the least-correlated members), N, and C (every crossing acquires an administrative cost). *Signature:* crossing counts drop far more for junior and new members than senior ones; cross-team encounters fall while within-team hold; someone is employed to operate the threshold. *Fix:* governance, not layout. **One gate at the perimeter; none inside.**

THE OPEN PLAN WITHOUT IMPEDANCE. All internal thresholds removed in the name of collaboration; M = 0 everywhere. *Damages* C catastrophically (k unbounded), u (Chapter 14's reverberant-room mechanism: the loudest voice wins by physics), and eta, which falls paradoxically because deep work produces nothing durable. *Signature:* k above 20/day; widespread headphone use — **headphones are occupants building, at their own expense, the threshold you refused to build**; booking pressure on the enclosed rooms; deep work relocated to homes and evenings. *Fix:* reintroduce impedance per S15.9, closable boundaries per S15.10, paired per S15.11. Do not revert to cellular offices; that overshoots, destroying the edge and the incidental encounter along with the noise.

THE AIRLOCK THAT IS A WAITING ROOM. A transition zone given a secondary function requiring dwell — queue seating, printing, coats — so it is permanently occupied and transitions nobody. Its sibling is the vestibule at the wrong depth: under 1.5 m and people stop hard or arrive at speed; over about 8 m and they treat it as corridor and arrive at speed anyway. *Damages* C in both cases — the transition fails and the mismatch lands at the next junction. *Fix:* separate the dwell function out, or move it one zone outward; add a bend to the shallow vestibule and subdivide the deep one with a marked step.

THE INVISIBLE THRESHOLD. A zone change existing in the organisation's mind with no physical marker: the "quiet end", the "team area". *Damages* C (enforcement is social, so every violation costs a human interaction) and u (the burden falls on whoever will speak up, who are not a random sample). *Signature:* repeated verbal enforcement of a spatial norm; a printed sign asking for quiet — **signs are the receipt for a failed geometry**. *Fix:* make it physical per S15.5. **The test is whether a person never told the rule obeys it on their first crossing.**

THE UNMEASURED SWITCH. A threshold intervention justified, and later evaluated, on how people say the space feels. *Damages* not a parameter of the law but a parameter of your epistemics. *Signature:* the post-occupancy report has satisfaction scores and no counts; nobody can state k or m for either condition. *Fix:* S15.13 — and if you have already touched it, say so, measure now, and label the result a level, not a change.

Worked Example

Step 1 — The system and its parameters

A 30-person research floor, 900 m² net: single open plate 22 m x 41 m, partial glazed perimeter, four small enclosed rooms at the core, one entrance door from the lift lobby straight onto the working floor. Nine specialisms; four years old; holding a real stock of shared method, tooling and working relationship — its Q.

Surveyed per Chapter 24, with Paper E §5.1's operational choices stated:

- **Order observable `Q`** — normalised count of live shared methods, tools, protocols and documented working agreements in use. Measured `Q_0 = 6.0`.
- **Interaction observable `U`** — disagreement audit across the nine specialisms, `u = 0.75`.
- **Categories `N = 9`**, verified non-redundant by Chapter 7's collision test.
- **Conversion `eta = 0.45`**, durable structure per unit surplus over eight quarters.
- **Decay `delta = 0.06/week`**, from decay under zero input across two quiet periods.
- **Cost observable `C` — calibration choice, declared openly:** `C = phi * eta * ln N`, `phi` a dimensionless fraction of the ceiling. Paper E §5.1 requires the cost observable to be chosen and declared; this is the declaration, not a derivation.

```
ln N = ln 9 = 2.1972    N_eff = 1 + 8(0.75) = 7.000    S = ln 7 = 1.9459
eta*lnN = 0.9888  <- ceiling   eta*S = 0.8757  <- input   delta*Q = 0.3600  <-
maintenance
```

The floor runs at 88.6% of its geometric ceiling. **There is almost no headroom left to buy by improving diversity; Chapter 9's levers are nearly exhausted here.** The problem is elsewhere.

Step 2 — Measure C

Protocol 15.4 (tally cards, two weeks, 30 people): median `k = 22`. Protocol 15.3 (progress-rate method, twelve person-days): median `m = 7 minutes`, interquartile range 3 to 16.

```
C_switch  = 22 * 7 = 154 min/day of 480 = 0.321
phi_other = 0.25     (meeting overhead, tooling friction, compliance, room booking)
phi_before = 0.571   ->  C_before = 0.571 * 0.9888 = 0.5646
```

Step 3 — Headroom before

```
h = eta ln N - delta Q - C = 0.9888 - 0.3600 - 0.5646 = +0.0642
```

Headroom of +0.064 against a ceiling of 0.989 — six per cent of the maximum one-step increment. By Corollary 4.2, `r_max = h`: any project asking for more growth per week cannot be delivered, whatever is promised. Within measurement error the floor is at maintenance only (Paper E §6.4: hold, do not build).

Steady state, Theorem 3.5: `Q_inf = (0.8757 - 0.5646)/0.06 = 5.185`. Current `Q_0 = 6.0` is **above** its own steady state: the floor is not stagnant, it is declining slowly toward 5.19, everyone can feel it, and nobody can name a cause — because the cause is distributed across twenty-two interruptions a day.

Step 4 — Diagnosis

Mismatch at the entrance. Lift lobby (public, transit, strangers, raised) to working floor (common, task, low, deep): `M = 4`, `M_bar = 2.4`. **A cliff edge.**

Edge density. Enclosing perimeter `2(22 + 41) = 126 m` on 902 m² = **0.140 m/m²** against a required 0.40; core rooms add about 44 m, half occupiable — effective **0.11 m/m²**. A severe deep-plan condition: perimeter desks occupied 94% of working hours, interior desks 38%.

Desire-line survey, six weeks. A diagonal from the entrance across the floor to the kitchen, cutting through two clusters; chairs migrating toward the glazing week on week; three of four core-room

doors habitually propped; handle wear on the fourth an order of magnitude above the rest, which is where the private conversation happens. Four closable boundaries shared by 30; no screens; headphones in use by 21 of 30 — twenty-one people privately funding the threshold layer the building did not supply.

Step 5 — Intervention

Chosen against S15.2, S15.6, S15.9, S15.10 and S15.11. Nothing changes the roster, recruitment, meeting culture or management structure. It is all joinery.

1. **Entrance sequence.** 30 m² beside the lift lobby becomes a vestibule (2.4 x 2.4 m, surface change, lower ceiling, warmer and lower light) and an anteroom (2.4 x 3.0 m, two seats, sightlines both ways), entry turned 90deg. `M` at every junction falls from 4 to 1 or 2.
2. **Re-route the kitchen path.** Kitchen moved beside the anteroom; the diagonal becomes a short public-zone route crossing no cluster. *This one move removes the largest k source.*
3. **Manufacture edge.** Fourteen thin islands (0.9 x 2.4 m storage, 1.4 m high per S15.9) form eight semi-private clusters of three or four desks at 0.9 m and 1.2 m clearances. Edge added `14 x 2 x (0.9 + 2.4) = 92.4 m`; new density `(126 + 44 + 92.4)/902 = 0.291 m/m²`; floor consumed 30.2 m² (3.3%).
4. **Corridor pockets.** Three widenings of 1.2 x 2.4 m on the main route, each with a leaning rail and a writable surface — 8.6 m².
5. **Closable boundaries.** Two core rooms become bookable focus rooms with working locks; the propped pair get alarm-released hold-open magnets.
6. **Paired invitation (S15.11).** Six window places at the glazing: unassigned, first-come, no booking, no ownership, open to visitors. The counterweight to items 3 and 5, and not optional.
7. **Availability protocol.** Two protected blocks daily, 09:30–11:30 and 14:00–15:30. Protocol, not architecture — subject to δ, reviewed quarterly forever.

Total area consumed **68.8 m², 7.6% of net floor.**

Step 6 — Parameters after

Re-measured at eight weeks, same protocols, same observers:

```
k = 6 per person per day        (was 22)
m = 7 minutes             (unchanged, as expected: m is a task property, not a spatial one)

C_switch  = 42/480 = 0.087
phi_after = 0.087 + 0.25 = 0.337   ->  C_after = 0.3332
h_after   = 0.9888 - 0.3600 - 0.3332 = +0.2955
Q_inf     = (0.8757 - 0.3332)/0.06  = 9.041      (was 5.185)
```

`eta`, `delta`, `N` and `u` are taken as unchanged, deliberately: **the intervention touched exactly one term of the law.** Headroom rises 4.6-fold and the sustainable level of stored order by **74%**, with no change to the people, the work, or the diversity of the group.

Step 7 — The transient

From Theorem 3.1, `Q_t = (1-delta)^t Q_0 + (1 - (1-delta)^t) Q_inf`:

t (weeks)	0	5	10	20	30	50
Q_t	6.000	6.809	7.403	8.159	8.566	8.903

Time constant `tau = 1/delta = 16.7 weeks`; 95% settling at about `3 tau = 50 weeks`. **Do not evaluate this intervention at eight weeks on the order observable.** At eight weeks the floor has moved to about 7.2 — 31% of the total — and reading that as a disappointment misreads a first-order step response. Evaluate `C` at eight weeks, because `C` is the input; evaluate `Q` at a year. **The cost term responds immediately and the stored order responds at the time constant, so the two must be measured on different schedules or the intervention will be judged a failure while it is succeeding.**

Step 8 — What would falsify this

Stated in advance, per Chapter 25:

1. **k does not fall.** k above 15 at eight weeks means the threshold layer is not filtering and the model is wrong. Log interruption *origin* as well as count, to separate the kitchen re-route from the islands.
2. **k falls but `Q` does not.** `Q_20` below 7.4 means recovered time is not becoming durable structure, `eta` is lower than measured, and `C` was never binding. Stop; go to Chapter 4.
3. **u falls.** A six-month disagreement audit below 0.65 means the clusters have become silos and the thresholds have partitioned the group's epistemics along with its noise — Chapter 20's failure mode installed by this chapter's own tools. Rotate cluster membership quarterly and check that the window places serve more than one cluster. **If one cluster has colonised them, S15.11 has failed and the filter is doing damage.**
4. **The shuffled-input control.** A matched floor, same organisation and period, given a cosmetic refurbishment of equal cost and disruption but no change to `M`, k or edge density. If its `Q` rises comparably, the claimed mechanism is not the operating one. **The control most often omitted, and the one that decides whether this is engineering or decoration.**

The honest ledger. Demonstrated: a cost-side intervention changing no personnel, no diversity and no conversion can multiply headroom fourfold and raise sustainable stored order by three quarters, for under 8% of floor area. Not demonstrated: that these are the numbers you will find. **Every figure here is an instance, not a constant.** The method transfers; the numbers do not.

And one thing did not change: **S15.6 was not met.** Edge density reached 0.291 m/m² against a requirement of 0.40, because the plate is 22 m deep and no arrangement of islands fixes a 22 m plate. The specification records the shortfall rather than relaxing to meet the design — a specification that adjusts itself to what was achieved has stopped being a specification. The correct project-record entry is: *S15.6 not met; deficit 0.109 m/m²; cause, plate depth; remedy requires subdivision or a different building.* That entry is worth more to the next engineer than any of the successes above.

Chapter 16 . Light, Material, and Thermal

This is the substrate layer. Chapters 12 to 15 assumed a room a person could stand to be in for a long time; this chapter is about what makes that true. The substrate is treated as decoration by almost everyone who specifies it and as physics by almost nobody — a category error with an arithmetic cost. Three claims:

1. **Space affects outcome, and there is real clinical evidence for it.** One small study from 1984 is still the cleanest demonstration, worth more than a thousand preference surveys because it measured a clinical result rather than an opinion.
2. **Thermal comfort has six variables, not one.** Almost every building controls one. The other five are where the complaints come from.
3. **Individual adjustability beats an optimal fixed setpoint** — not because optimisation is done badly, but because the optimum of a dispersed population is a place almost nobody is.

The Law

What is proved

From Paper E §3:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t          S_t = ln(1 + (N−1) U_t)
```

and the headroom gauge (Paper E §6.4):

```
h_t := eta ln N − delta Q_t − C_t
```

Nothing about light, air or material is derivable from these. The law supplies the accounting frame: once you have established that a substrate variable moves δ, η, C or N, the law prices the move. This chapter's work is establishing the *that*, and being explicit about the evidence.

From Paper D, Corollary 2.2, `0 <= F(u) <= ln N`. **Surplus is bounded above by the log of the number of genuinely distinct categories present.** Anything that removes a category from the room — including a body that cannot tolerate the room — lowers the ceiling itself, not merely the achieved value. §16.4 is about that mechanism.

What is postulated

POSTULATE 16.1 (Substrate). The physical substrate bounds the duration and quality of voluntary occupancy, and voluntary occupancy carries interaction. Substrate therefore enters the law through `eta` (occupied time is the substrate of durable relation), `C` (discomfort is dissipation, and so is every adjustment negotiation) and `delta` (a room that degrades needs replacing). **Badged as a modelling postulate**, supported by direct observation and §16.1's clinical evidence; not a theorem.

POSTULATE 16.2 (Adjustability). For a comfort variable over which individual preferences are dispersed, distributed individual control outperforms a single optimised fixed value, measured by the fraction of occupants outside their acceptable band. **Badged** — but with an arithmetic backbone, given in §16.4: the standard comfort model's own dissatisfaction floor is 5% at the

theoretical optimum. **The model used to justify fixed setpoints states in its own terms that a fixed setpoint cannot satisfy everyone.**

POSTULATE 16.3 (Patina). A material whose surface reads better after wear has a lower effective `delta` on the appearance component of stored order than one that only degrades. **Badged.** Jacobs (*Death and Life of Great American Cities*, 1961) argued for aged buildings on economic and social grounds; the δ framing is this manual's interpretation, not her finding.

Honesty clauses carried into this chapter

- Paper A is pure mathematics and makes **no empirical claim**. No result about daylight, temperature or CO2 follows from it. As deployed, the surplus functional is a **lower bound** on the ensemble's information, and **the triangle inequality fails for geometric triples (Paper A §5)** — do not add comfort "distances".
- Where this chapter cites a study it states the sample size and the design; where a literature is contested it says so in the same paragraph as the claim, not in a footnote. §16.5 is the test case, and it is deliberately unflattering to the argument this chapter would prefer to make.
- Numbers taken from standards are engineering-practice values in wide use, and they change between editions and jurisdictions. **Get the current edition of the standard that governs your project. This is a manual about method, not a substitute for the code.**

The Engineering

16.1 Ulrich 1984: the strongest single piece of evidence, stated precisely

Roger Ulrich, "View Through a Window May Influence Recovery from Surgery," *Science*, 1984. The design: a retrospective study of records from a suburban Pennsylvania hospital, 1972–1981. Patients recovering from cholecystectomy (gall-bladder removal) were matched in pairs on age, sex, smoking status, weight, previous hospitalisation, year of surgery and attending surgeon. **Twenty-three matched pairs.** The rooms were essentially identical except for what the window faced: a stand of deciduous trees, or a brown brick wall.

The findings: tree-view patients had shorter postoperative stays, received fewer doses of moderate and strong analgesics, and drew fewer negative evaluative comments in the nurses' notes ("upset and crying," "needs much encouragement"). The stay difference was a fraction of a day against a mean of roughly eight.

Why this is worth more than the surveys. It measured *outcomes* — length of stay, analgesic dose, clinical notes — not preferences. Nobody was asked whether they liked the view, the patients did not know they were in a study, and the outcome measures came out of the hospital's ordinary operation. That removes the demand characteristics that make almost all built-environment preference research uninterpretable.

Why to be careful, said as loudly as the finding. Twenty-three pairs is a small sample. One hospital, one procedure, one climate, one decade. It was retrospective, and room assignment was as-available rather than randomised — close to random in practice, not guaranteed. The effects are modest, and the nurses' notes were scored from records not written for the purpose. **If you intend to quote the numbers, read the paper and take the figures from it, not from this or any other secondary source.**

The subsequent literature on nature exposure and health is large, heterogeneous and mixed: some effects replicate, some do not, sizes vary widely, much of it is cross-sectional. **The honest summary is that the direction is consistent, the magnitude is uncertain, and Ulrich remains unusually clean rather than unusually large.** The engineering conclusion is not "views cure people." It is narrower and more useful:

A variable the profession classifies as aesthetic turned up in a clinical outcome measure, so the classification is wrong. Light, view, material and air are an input layer, and must be specified, measured and defended like any other. That licenses the rest of this chapter. It does not license inventing effect sizes.

16.2 Light

Illuminance, and the plane you measure it in Standard maintained illuminance ranges in current lighting practice: **100–150 lux** circulation and corridors; **300–500 lux** general office, reading and meeting; **750–1000 lux** detailed work, drawing and inspection; **1000–1500 lux** very fine and colour-critical work.

For scale: overcast sky delivers roughly 5,000 to 20,000 lux outdoors, direct sun 50,000 to 100,000, full moonlight under 1 lux. **A well-lit office delivers about 3% of a cloudy day and 0.5% of full sun.** Everyone who walks outside at lunch has felt that ratio without naming it.

Colour: CRI Ra of 80 or better generally, 90 for colour-critical work. Correlated colour temperature is a design choice, not a comfort constant — 2700–3000 K domestic, 4000 K neutral, 5000–6500 K daylight-like.

Glare is controlled by limiting unified glare rating (UGR <= 19 general office, <= 16 drawing, <= 25 corridors) and by limiting luminance ratios: roughly 3:1 task to immediate surround, 10:1 task to remote surfaces. **The commonest glare source in a well-designed room is the window** — the central tension of this section, since the best circadian source and the worst glare source are the same object.

Daylight Daylight factor is internal over unobstructed external diffuse illuminance.

DF	5,000 lux out	10,000	20,000	Character
0.5%	25 lux	50	100	electric light always needed
1%	50	100	200	poorly daylit
2%	100	200	400	conventionally "daylit"
5%	250	500	1000	well daylit, glare risk

The rule that decides the plan: **useful daylit depth from a side window runs roughly 2.0 to 2.5 times the window head height** — 4.2–5.2 m at a 2.1 m head, 4.8–6.0 m at 2.4 m, 5.4–6.8 m at 2.7 m, 6.0–7.5 m at 3.0 m.

A rule of thumb, dependent on glazing transmittance, external obstruction, internal reflectance and orientation — and nonetheless the number that kills more plans than any other. A room daylit from one side cannot be much more than 6 m deep; daylit from two sides (Alexander's Pattern 159, *Light on Two Sides of Every Room*) it can be about 12 m.

This is Chapter 15's edge rule arriving through the window. A 22 m plate cannot reach edge density 0.40 m/m² and cannot be daylit past its perimeter strip, and both failures have the same cause: the ratio of perimeter to area. A plan that violates one violates the other.

Circadian entrainment The eye contains intrinsically photosensitive retinal ganglion cells driving the circadian system, with a spectral sensitivity peaking around 480 nm, differing from the photopic curve the lux unit is built on. The quantity that matters is therefore not lux but **melanopic equivalent daylight illuminance (mEDI)**, measured at the eye.

An expert-consensus recommendation published in 2022 (Brown and colleagues, *PLOS Biology*) proposes at least **250 lux mEDI** at the eye during the daytime, under **10 lux mEDI** in the three hours before bed, and under **1 lux mEDI** during sleep. **Badge: a consensus recommendation from a panel, not a regulation and not a settled constant; the literature is active and the thresholds may move.** It is still the best available target, and far better than the tacit alternative of none.

The lux meter on the desk is measuring the wrong plane. Task illuminance is horizontal; the circadian system responds to light at the eye, a vertical plane facing the direction of gaze. A 500 lux desk under downlights can deliver under 150 lux vertical at the eye. **Measure vertically, at seated eye height, facing the normal direction of view**, and report both planes. A room can pass the task standard and fail the circadian target threefold, and every current lighting sign-off will call it compliant.

Second, the evening half of the recommendation specifies what to *turn off*. A building delivering 250 mEDI at 15:00 and again at 22:00 has satisfied one half and violated the other. Time-varying output is the specification, not a luxury.

The window is both Summarising the tension: daylight has very high melanopic content, free and correctly phased time variation, and the view Ulrich measured; it also has high glare risk, poor uniformity, and it is unavailable at night and in a deep plan. Electric light is always available, uniform and glare-controllable at design stage, but its melanopic content is low unless deliberately designed for, and its time variation costs a control system.

Design rule 16.A. Daylight the perimeter and control its glare at source; supplement the interior with a scheme designed for vertical eye illuminance rather than horizontal desk lux; vary it over the day. Do not shade the windows until the daylight advantage is thrown away — the commonest outcome of a glare complaint is a permanently closed blind, which converts a solvable glare problem into a permanent light and view deficit.

16.3 Thermal comfort: the six variables

Six variables — four environmental, two personal:

#	Variable	Symbol	Units	Typical indoor range
1	Air temperature	t_a	degC	19 – 27
2	Mean radiant temperature	t_r	degC	15 – 28
3	Relative humidity	RH	%	25 – 65
4	Air speed	v	m/s	0.05 – 1.0

#	Variable	Symbol	Units	Typical indoor range
5	Clothing insulation	I_cl	clo	0.3 – 1.2
6	Metabolic rate	M	met	0.8 – 2.0

This is the set used by the standard comfort models (Fanger's PMV/PPD, as embodied in ISO 7730 and ASHRAE Standard 55). **Almost every building controls variable 1 and ignores the other five** — the source of most thermal complaint, and the reason turning the thermostat does not fix it.

```
CLOTHING clo (1 clo = 0.155 m2K/W)        METABOLIC RATE met (1 met = 58.2 W/m2)
  nude                     0.0                    sleeping                  0.7
  shorts + T-shirt         0.36                   seated, quiet             1.0
  typical summer office    0.5                    seated office work        1.2
  trousers + shirt         0.61                   standing, light work      1.6
  + jacket                 0.96                   walking 3 km/h            2.0
  typical winter office    1.0
```

Operative temperature — the single number that predicts comfort at low air speed — is approximately the mean of air and mean radiant temperature:

```
t_op ~= (t_a + t_r) / 2          (valid for air speed below about 0.2 m/s)
```

Mean radiant temperature is the invisible variable `t_r` is the area-weighted mean of surrounding surface temperatures seen from the occupant's position. Nobody measures it; it causes much unexplained discomfort. Glazing inner surface temperature:

```
T_si = T_i - U (T_i - T_o) / h_i        h_i ~= 7.7 W/m2K (internal surface coefficient)
```

With the room at 21degC:

Glazing (U, W/m²K)	Outside -5degC	0degC	+5degC
Single (5.7)	**1.8degC**	5.5degC	9.2degC
Double, air (2.8)	11.5degC	13.4degC	15.2degC
Double low-E argon (1.4)	16.3degC	17.2degC	18.1degC
Triple low-E (0.8)	18.3degC	18.8degC	19.3degC

For someone close to the glass, roughly half the field of view is window, so the *local* mean radiant temperature is about the average of the glazing surface and a room at 21degC:

Glazing inner surface	8degC	12degC	15degC	17degC
Local t_r	14.5degC	16.5degC	18.0degC	19.0degC
Local operative T (air 21degC)	17.8degC	18.8degC	19.5degC	20.0degC
Deficit	**3.2 K**	2.2 K	1.5 K	1.0 K

The person by the single-glazed window in winter is sitting in an effective 17.8degC while the thermostat truthfully reports 21degC. They complain; facilities raises the air to 23degC, lifting their operative temperature to 18.8degC — still cold — while the middle of the room is now too hot. The complaint has been answered in the wrong variable, twice, at a permanent energy cost.

The fix is a fabric fix: better glazing, a radiant surface, or a different seat. **Chapter 15's S15.8 requires a seated position within 1.5 m of the glazing, because Ulrich says the view is worth having. If that seat sits behind single glazing in January, the view is being paid for in 3.2 K and nobody will use it.** The light chapter and the thermal chapter must be specified together or they cancel.

Note that the whole-room area-weighted MRT barely moves: a 17.4 m² glazed wall in a 159 m² envelope shifts room-average operative temperature only from 20.9degC to 20.3degC as the glass goes from 19degC to 8degC. **The room average is fine. The seat is not. Measure at the seat.**

Humidity and air speed Relative humidity has weak effects on thermal sensation within 30–60% and strong effects outside it: below 20–30%, dry eyes, dry mucosa and static; above 60%, mould risk, dust mites and a sticky sensation at raised temperature. **Specify a band, not a setpoint, and specify it for the fabric as much as for the person.**

Air speed cuts both ways with a sharp threshold. Below roughly 0.15–0.20 m/s air is felt as still; above that, at neutral or cool temperatures, it is felt as draught — worse at neck and ankles, worse with turbulence, worse for anyone already cool. At warm temperatures the same movement is relief, and the standards permit elevated air speed to offset a raised operative temperature provided the occupant controls it. **The offset table is in ASHRAE 55; get it from the standard, because it depends on clothing and on how the control is provided.**

Occupant-controlled local air movement is the single best-value comfort device in the catalogue. A desk fan costs less than an hour of an engineer's time, is instantly reversible, and hands the sixth variable to the person who owns the body.

16.4 Why adjustability beats the optimum

The arithmetic The predicted mean vote (PMV) runs from -3 (cold) to +3 (hot), and the predicted percentage dissatisfied is derived from it:

```
PPD = 100 - 95 * exp( -(0.03353 PMV^4 + 0.2179 PMV^2) )
```

Computed:

PMV	0.0	0.2	0.3	0.5	0.7	1.0	1.5	2.0
PPD	**5.0%**	5.8%	6.9%	10.2%	15.3%	26.1%	50.9%	76.8%

Read the first column again. **At the theoretical optimum, every variable ideal, the model predicts 5% dissatisfied. The floor is not zero.** The usual criterion, |PMV| <= 0.5 giving PPD <= 10%, accepts one person in ten uncomfortable *by design* — in a building working exactly as specified.

That floor is not a defect of the model. It is the model correctly reporting that thermal preference is dispersed and that one environmental state cannot sit at every individual optimum. **A fixed setpoint, however well optimised, answers the question "where is the centre of the distribution?" — and almost nobody lives at the centre of a distribution.**

What adjustability buys Model individual preferred operative temperatures as normal with standard deviation `sigma`, and suppose each person accepts ±1 K around their own preference. **`sigma` is an assumed parameter, badged as such; measure it in your own population (Protocol 16.1).** The 1.5–2.5 K range is the plausible design region. Fraction satisfied by one fixed setpoint at the population mean:

σ (K)	±1.0 K band	±1.5 K	±2.0 K
1.0	68.3%	86.6%	95.4%
1.5	49.5%	68.3%	81.8%
2.0	38.3%	54.7%	68.3%
2.5	31.1%	45.1%	57.6%

Give each occupant personal authority of ±A K over their own operative temperature, widening the effective acceptance band to ±(1 + A):

Personal authority A	σ = 1.5 K	σ = 2.0 K
±0 K (fixed setpoint)	49.5%	38.3%
±1 K	81.8%	68.3%
±2 K	95.4%	86.6%
±3 K	99.2%	95.4%

At σ = 2.0 K, two degrees of personal authority takes satisfaction from 38% to 87%. No amount of re-optimising the single setpoint can do that, because the problem is the dispersion, not the location. That is the whole argument, and it is arithmetic rather than sentiment.

It is robust to the assumed σ in the way that matters: the *ratio* of adjustable to fixed performance improves as σ grows. A population more varied in age, body composition, medication, metabolic rate, gender or dress makes the case stronger, not weaker.

The four adjustment channels, ranked by cost

Channel	Variable moved	Cost	Speed	Personal or shared?
Clothing	I_cl	zero	seconds	**personal**
Local air movement (fan)	v	trivial	seconds	personal
Local radiant (small panel, warm surface)	t_r	low	minutes	personal

Channel	Variable moved	Cost	Speed	Personal or shared?
Zone setpoint	t_a	high	tens of minutes	**shared — and therefore contested**

Clothing is the most powerful and cheapest lever and the one organisations routinely disable. As a rule of thumb from the PMV model, 0.1 clo shifts the neutral operative temperature by roughly 0.6 K, so moving between a 0.5 clo summer outfit and a 1.0 clo winter one is worth about 3 K. **Badge: verify against the ISO 7730 tables for your metabolic rate before designing to it.**

A dress code is thermal policy. Fixing clo removes the free, personal, instantaneous channel and transfers the whole load onto a shared, slow, expensive one, so the building burns energy to compensate for a decision taken in a handbook. If occupants must wear jackets in July, the cooling bill is a clothing policy expressed in kilowatt-hours.

The cost of contested adjustment Zone setpoint control is *shared*, so every adjustment is a negotiation. It has a name in every office in the world and it belongs in `C_t`.

Instrument. Log thermostat interactions — who, when, direction, magnitude, and whether a conversation followed. A shared control changed more than about twice a day is not a control, it is a dispute. Convert to the cost observable as contested adjustments per period x mean time consumed per adjustment, including the discussion and the counter-adjustment. Small in isolation, persistent, and — unlike most C terms — trivially measurable, because thermostats can log.

Design rule 16.B — never install a shared control where a personal one will do. A personal control changes one person's environment and generates no coordination cost; a shared control generates one on every use, forever, and it falls hardest on whoever is least willing to argue.

The connection to u, made carefully The claim the chapter wants to make, with its badge on:

POSTULATE 16.2a. A room that can only be set one way selects for the bodies it suits. People whose thermal preference lies far from the setpoint attend less, disengage sooner, relocate or leave. Since thermal preference is not independent of age, body composition, gender, health status or medication, that selection is not random with respect to the categories the surplus law counts. **A fixed setpoint therefore reduces `N_eff` by attrition, and lowers `ln N` itself if a category is lost entirely.**

This is the most speculative claim in the chapter. It is stated because it is testable.

Protocol 16.1 — the thermal attrition audit. Collect individual preferred operative temperature by structured self-report (a seven-point sensation scale under measured conditions, at least ten occasions per person across a season). Rank occupants by distance from the delivered operative temperature, then measure by decile: hours present, voluntary attendance at optional gatherings held there, and twelve-month retention. **If distance from setpoint predicts none of the three, Postulate 16.2a is false in your building and should be dropped.** Report either result.

The well-supported half of the story needs no postulate: adaptive comfort. Field studies in naturally ventilated buildings — the de Dear and Brager work behind the adaptive model now in ASHRAE 55 — find that occupants with real control accept a substantially wider range of conditions than the laboratory PMV model predicts. **Perceived control widens the acceptance band independently of what it is used to do.** That is replicated, and on its own it justifies Design rule 16.B.

The honest costs of adjustability

- **Energy.** Wider zone tolerance plus personal devices is usually a net saving, but personal resistance heaters are the ruinous exception. Specify personal *radiant* and *air movement* devices, not personal fan heaters.
- **Control complexity.** Every controllable device is a maintenance item. δ applies to control systems more brutally than to fabric.
- **Spillover.** The channel must be personal in its *effect*, not merely in its interface. A desk fan blowing across a neighbour's neck has converted a personal control into a shared one.

16.5 Air quality, and the CO2 literature stated honestly

CO2 as a tracer Steady-state indoor CO2 under constant occupancy is

```
C_in = C_out + G / Q
  C_out  outdoor concentration, about 420 ppm and rising
  G      CO2 generation per person, about 0.0042 to 0.0060 L/s at 1.2 met
  Q      outdoor air supply per person, L/s
```

At `G = 0.005 L/s` per person:

Outdoor air Q (L/s per person)	2.5	4	5	8	10	15	20
Rise above outdoor (ppm)	2000	1250	1000	625	500	333	250
Steady-state C_in (ppm)	2420	1670	1420	1045	920	753	670

Inverted, so you can read ventilation off a logger: a measured steady-state 600 ppm implies 27.8 L/s per person; 800 ppm implies 13.2; 1,000 implies 8.6; 1,200 implies 6.4; 1,500 implies 4.6; 2,000 implies 3.2.

This inversion is the reason to own a CO2 logger: it turns a cheap instrument into a measurement of outdoor air per person, otherwise expensive to determine and almost never verified after commissioning.

Now the uncomfortable arithmetic. ASHRAE 62.1's office breathing-zone rate is 2.5 L/s per person plus 0.3 L/s per m²; for a 42 m² room with 4 people:

```
Q = 2.5 x 4 + 0.3 x 42 = 22.6 L/s total = 5.65 L/s per person
steady-state rise = 0.005 x 4 / 0.0226 = 885 ppm  ->  C_in ~= 1,305 ppm
```

A room ventilated exactly to standard settles at about 1,300 ppm — above the concentrations at which the cognitive-performance studies reported effects. Either the standard is wrong, or those studies are wrong, or the effect does not matter. This chapter cannot resolve that, and says so.

The literature, stated honestly The claim in circulation is that CO2 at concentrations common in occupied rooms — roughly 1,000–2,500 ppm — impairs decision-making. Its main sources are Satish and colleagues (2012, *Environmental Health Perspectives*), reporting large decrements on a computerised decision-making instrument at 1,000 and 2,500 ppm, and Allen and colleagues (2016, same journal), reporting similar effects in a controlled office simulation.

The honest state of that literature, which anyone quoting it is obliged to state:

- **It has had replication difficulties.** Several subsequent controlled studies failed to reproduce the effect; work with submarine-relevant protocols, at concentrations far above office levels, reported no decrement on cognitive test batteries. Reviews of the field consistently describe the evidence as inconsistent.
- **Samples are small** and outcome instruments heterogeneous, so between-study comparison is weak.
- **Blinding is hard.** Elevated CO2 in a real room comes with other changes participants can sometimes detect.
- **The central confound is structural.** In any real building CO2 rises *because* ventilation is low, so bioeffluents, VOCs, humidity and thermal load rise with it. **Dosing pure CO2 into a well-ventilated room tests something different from what a stuffy meeting room does, and the two designs disagree with each other.**

Conclusion, deliberately unhelpful to the case this chapter would like to make: the proposition that CO2 itself impairs cognition at office concentrations is contested, has had replication difficulties, and should not justify a ventilation investment. If you have quoted it in a business case, withdraw the claim and re-argue on the grounds below, which are stronger.

What is well supported

- Low ventilation rates associate with higher prevalence of sick-building symptoms — headache, fatigue, eye and airway irritation — across a substantial occupational-health literature.
- Rebreathed-air fraction governs airborne infection transmission risk, and CO2 is a direct proxy for it. A physical relationship, not a psychological one.
- Perceived air quality and odour acceptability improve monotonically with ventilation rate.

Engineering rule 16.C: specify ventilation, monitor CO2, and claim only ventilation. Use concentration as a tracer for a rate; claim no cognitive effect.

The rest of air quality

- **Particulate.** PM2.5 has an outcome literature vastly stronger than CO2's. Filter to ISO ePM1 or MERV 13 equivalent where outdoor PM is a concern, and verify the fan can drive the pressure drop — a good filter in an under-powered unit gets removed at the first service visit.
- **VOCs.** Control at source with low-emission finishes, adhesives and furniture; dilution is the expensive way to solve an emission problem. Ventilate hard during and after fit-out.
- **Instrumentation.** NDIR CO2 sensors with a documented calibration interval, at breathing height, at least 1.5 m from any face, away from doors and diffusers. Log continuously at 1–5 minute intervals; spot readings are near worthless, because the quantity is a slow integrator.

16.6 Materials

Materials do four jobs. Three are conventionally recognised; the fourth is this chapter's contribution.

Acoustic Absorption coefficients at 500 Hz, standard practice values:

Surface	α	Surface	α
Painted concrete, hard plaster	0.02–0.05	Heavy drapes, deep folds	0.50–0.60
Glass, glazed partition	0.03–0.05	Mineral-fibre ceiling tile	0.60–0.80
Gypsum board on studs	0.05–0.10	50 mm mineral wool panel	0.80–0.95
Carpet tile on concrete	0.15–0.30	Occupant, upholstered seat	~0.45 m² Sab each

Sabine, from Chapter 14: `RT60 = 0.161 V / A`, `A` the sum of area x coefficient. Worked for the 42 m² room used later (6.0 x 7.0 x 2.9 m, V = 121.8 m³, S = 159.4 m²):

```
HARD CASE   exposed concrete soffit 42 x 0.02  =  0.84
            plastered walls        67.4 x 0.03 =  2.02
            glazing                 8.0 x 0.04 =  0.32
            vinyl on concrete      42.0 x 0.03 =  1.26
            A = 4.44 m2Sab      ->  RT60 = 0.161 x 121.8 / 4.44 = 4.41 s
```

Four and a half seconds. That is a church, not a workroom, and it is what an all-hard room of this size genuinely does. Speech is unintelligible at conversational level, and by Chapter 14's mechanism the loudest voice wins by physics.

Thermal mass Volumetric heat capacity and diurnal penetration depth `d = sqrt(alpha P / pi)`, with `alpha = k/(rho c)` and `P = 86,400 s`:

Material	ρ	c	k	ρc (MJ/m³K)	24-h depth
Dense concrete	2300	880	1.7	2.02	152 mm
Brick	1700	800	0.72	1.36	121 mm

Material	ρ	c	k	ρc (MJ/m³K)	24-h depth
Gypsum board	800	1000	0.25	0.80	93 mm
Solid timber	500	1600	0.13	0.80	67 mm
Water	1000	4180	0.6	4.18	63 mm

Penetration depth is where the daily swing has fallen to 1/e of its surface amplitude; standard design practice takes roughly the first **100 mm** of exposed heavyweight material as the effective participating depth. **A 300 mm slab exposed on one face therefore performs like a 100–150 mm slab. Thickening buys nothing for the diurnal cycle; exposing more *area* buys everything.**

Stored energy in exposed concrete:

```
per m2 of exposed surface, over the participating depth: 2300 x 880 x 0.152 = 308 kJ/m2/K
100 m2 of exposed soffit over a 3 K diurnal swing: 308 x 100 x 3 = 92.3 MJ = 25.6 kWh
```

Twenty-six kilowatt-hours of free storage, shifted from afternoon peak to night, from a soffit you already paid for.

And then: carpet the slab, or hang a sealed suspended ceiling under it, and the mass is gone. The material is still there; its coupling to the room air is not. **A decision taken for acoustics silently deletes the thermal strategy, and nobody notices, because the two decisions are made by different people in different months.**

The resolution is the treated case, computed:

```
TREATED   exposed soffit, 60% left open       25.0 x 0.02 =  0.50
          50 mm mineral wool baffles          17.0 x 0.90 = 15.30
          plastered walls                     57.4 x 0.03 =  1.72
          10 m2 wall panel, 50 mm             10.0 x 0.85 =  8.50
          glazing                              8.0 x 0.04 =  0.32
          carpet tile on concrete             42.0 x 0.25 = 10.50
          4 occupants                          4   x 0.45 =  1.80
         A = 38.64 m2Sab   ->  RT60 = 0.51 s      exposed mass retained: 60%
```

RT60 from 4.41 s to **0.51 s**, inside the speech range, with 60% of the soffit still thermally coupled. **Baffles and rafts absorb without sealing.** The two requirements are compatible; they are never coordinated.

Tactile Two surfaces at identical temperature feel different, and the physics is exact. Thermal effusivity `e = sqrt(k rho c)` governs the interface temperature reached when skin touches a surface:

```
T_contact = (e_skin T_skin + e_mat T_mat) / (e_skin + e_mat)
```

With skin at 33degC and effusivity taken as 1200, and every material at a true **20degC**:

Material	e	Contact T	Material	e	Contact T
Aluminium	22,045	**20.7degC**	Glass	1,369	26.1degC

Material	e	Contact T	Material	e	Contact T
Mild steel	13,539	21.1degC	Oak	436	29.5degC
Stainless steel	7,697	21.8degC	Pine	322	30.2degC
Granite	2,421	24.3degC	Wool carpet	125	31.8degC
Dense concrete	1,855	25.1degC	Cork	104	**32.0degC**

At the same 20degC, an aluminium handrail presents 20.7degC to the hand and a cork one presents 32.0degC — eleven degrees at the skin from a material choice alone. No thermostat setting produces that difference and none compensates for it.

Specify by effusivity wherever skin meets building: handrails, door handles and push plates, seat surfaces and arms, worktops, sills people lean on, floors in barefoot spaces. **The steel handrail on an unheated stair is a thermal decision being made without the calculation.**

Temporal: patina and δ Materials divide by what wear does to them:

Improve with wear	Degrade with wear
Oiled and waxed hardwood	Melamine and laminate faces
Brass, bronze, copper	Thin veneer on MDF
Vegetable-tanned leather	Printed vinyl and photo-laminate
Lime plaster, brick, stone	Painted MDF edges
Linoleum, cork	Powder-coat over thin steel

The distinction is structural, not aesthetic. **A material improves with wear when its surface and its body are the same substance**, so wear exposes more of the same. It degrades when a thin surface layer differs from the substrate, so wear breaches the layer and exposes something visibly other.

In the law's terms, `delta` is per-component. Structural δ (the desk stops being a desk) is small for both classes; appearance δ (the desk stops reading as cared-for) is large for layered materials and near zero — arguably negative over a decade — for solid ones. Since perceived maintenance drives the behaviour that actually maintains a place, **a low-appearance-δ material is a self-reinforcing investment and a high-δ one is self-defeating.**

Illustrative life-cycle arithmetic over 40 years, per m² of worktop:

```
laminate: 180 initial, replaced every 8 years -> 5 purchases = 900  (22.50 /m2/yr)
timber:  430 initial + 6/m2/yr oiling      -> 430 + 240  = 670  (16.75 /m2/yr)
```

Badge: those unit rates are illustrative placeholders showing the structure of the calculation. Substitute your own quoted rates — the conclusion depends entirely on them. What does not depend on them is the shape: a layered material has a step-function cost stream with disposal and disruption at each step; a solid one has a small continuous stream. Step functions are also what get deferred when budgets are tight, which is how a building arrives at year 12 looking abandoned.

Jacobs's argument about aged buildings was economic — old buildings house the low-margin, unproven, various uses new construction prices out. This is a different point in the same family: **material that ages well keeps a place legible as cared-for at low continuing cost, and legibility-as-cared-for is a large part of what makes people willing to be there.**

16.7 Mapping the substrate to the law

Substrate variable	Parameter	Mechanism	How to measure
Daylight and view	`eta`	Occupied hours at the perimeter; Ulrich	DF survey, vertical mEDI, occupancy by position
Circadian light timing	`delta`	Sleep -> next-day capacity -> skill retention	mEDI at eye height, morning and evening
Glare	`C`	Blinds closed, screens moved, complaints	Count of permanently closed blinds; UGR
Operative temperature at the seat	`C`, `N_eff`	Discomfort; contested adjustment; attrition	Globe thermometer at seat; Protocol 16.1
Personal control provision	`C`	Removes coordination cost per adjustment	Controls per occupant; shared-control log
Ventilation rate	`C`	Symptoms, infection risk, odour	CO2 logger inverted to Q; symptom survey
Acoustic absorption	`u`, `eta`	Turn-taking equality (Chapter 14)	RT60 by Sabine; speaking-share distribution
Exposed thermal mass	`C`	Peak-load shifting; free storage	Coupled surface area; internal temperature swing
Tactile effusivity	`eta`	Willingness to occupy and to touch	Material schedule; compute contact temperature
Patina behaviour	`delta`	Appearance decay rate	Annual fixed-position photographs

---

The Specification

S16.1 — Daylit depth and daylight factor. No regularly occupied workstation shall sit further than **2.5 x window head height** from its daylight source; where the plan exceeds this, provide a second aspect (Pattern 159) or record the shortfall as a named deficit. Target average daylight factor not less than **2%**; below 1%, the room is electrically lit at all times and shall be designed as such.

S16.2 — Illuminance and colour. Maintained horizontal illuminance per the applicable standard: **100–150 lux** circulation, **300–500 lux** general office and meeting, **750–1000 lux** detailed work. CRI Ra >= 80, >= 90 for colour-critical work.

S16.3 — Vertical measurement. Illuminance shall be recorded in **both planes**: horizontal at the working surface, vertical at seated eye height facing the normal direction of gaze. **A scheme verified only in the horizontal plane is not verified.**

S16.4 — Circadian target. Not less than **250 lux melanopic EDI** at the eye during daytime occupied hours, with the means to fall below **10 lux mEDI** in evening-occupied spaces. Badged as a consensus recommendation, not code; record it with its source and date.

S16.5 — Glare. UGR <= **19** general office, <= **16** drawing and inspection, <= **25** circulation. Provide operable shading at every window with usable partial positions — a shade with only "open" and "closed" will end up closed.

S16.6 — The six variables, measured at the seat. Any thermal comfort claim shall state all six with the values assumed: air temperature, mean radiant temperature, relative humidity, air speed, clothing insulation, metabolic rate. **A claim naming only air temperature shall be rejected as incomplete.** Operative temperature shall be measured at occupied seat positions with a globe thermometer, including the seat nearest the glazing — not at a wall thermostat, not at the room centre only.

S16.7 — Glazing and the window seat. Where S15.8 requires a seated position within 1.5 m of glazing, the glazing shall be specified so the local operative temperature deficit at that seat does not exceed **1.5 K** at the design outdoor condition. Compute `T_si = T_i − U(T_i−T_o)/h_i` and the local radiant effect; record both.

S16.8 — Personal control. Every regularly occupied position shall have at least **two** personal adjustment channels confined in effect to that position — local air movement and local radiant by default. Clothing counts only where no dress code constrains it.

S16.9 — Shared controls and humidity. Shared setpoint controls shall be logged (time, direction, magnitude); more than **two** contested adjustments per day on one control is a specification failure, not a behavioural one — subdivide the zone or add personal channels. Maintain **30–60% RH** in occupied hours, specified for the fabric and contents as well as the occupants.

S16.10 — Ventilation and CO2. Specify outdoor air rate per the applicable standard, then **verify it in use** by logging CO2 at 1–5 minute intervals and inverting `Q = G / (C_in − C_out)`. **Ventilation may be claimed. Cognitive performance may not** (§16.5).

S16.11 — Reverberation with mass retained. Meet the Chapter 14 RT60 target using baffles, rafts and vertical absorbers rather than sealed suspended ceilings, and record the percentage of heavyweight soffit left thermally coupled. **Not less than 50% shall remain coupled** where a thermal mass strategy is claimed.

S16.12 — Participating mass. Where thermal mass is claimed, compute participating depth `d = sqrt(alpha P / pi)` and state the coupled area. Mass behind a sealed ceiling, under carpet on a raised floor, or insulated on its room face counts as **zero**.

S16.13 — Contact surfaces. Schedule every surface routinely touched by skin — handrails, door furniture, seat arms, worktops, sills, barefoot floors — and compute contact temperature at the

design condition. In unheated or transitional spaces it shall be not less than **26degC** with skin at 33degC, ruling out bare metal.

S16.14 — Patina classification and appearance monitoring. Classify every visible material as improving-with-wear or degrading-with-wear by the surface-equals-body test; where design life exceeds 15 years, not less than **60%** of visible surface area shall be improving-with-wear. Photograph the same ten surfaces from taped camera positions annually — the only practical appearance-δ instrument, at one hour a year.

The Failure Modes

THE AIR-TEMPERATURE-ONLY BUILDING. One controlled variable of six. *Damages* `C` (persistent complaint and adjustment) and `N_eff` by Postulate 16.2a. *Signature:* complaints cluster by *position* rather than by person; the thermostat reads correct while the perimeter is cold; heating and cooling run at once in adjacent zones. *Fix:* measure operative temperature at the seats, fix the radiant problem in the fabric, add personal channels per S16.8.

THE COLD WINDOW SEAT. A seat provided at the glazing per S15.8, behind glazing that makes it 3 K cold in winter. *Damages* `eta` — the most valuable seat in the room is unoccupied for a third of the year. *Signature:* seasonal collapse in occupancy at the window position; personal heaters appear; the blind stays down for insulation and takes the view with it. *Fix:* S16.7. This is a glazing specification, not a heating one.

THE HORIZONTAL-ONLY LIGHTING SIGN-OFF. 500 lux on the desk, 150 lux at the eye. *Damages* `delta` via circadian disruption. *Signature:* the commissioning report contains one plane of measurement; occupants call the room gloomy while the meter says it complies. *Fix:* S16.3 — add vertical-plane sources rather than raising desk lux, which only adds glare.

THE PERMANENTLY CLOSED BLIND. A glare problem answered by an occupant with the only tool available. *Damages* `eta` (view and daylight both gone) and `delta` (circadian). *Signature:* count the blinds down on a dull morning; above 20% is a diagnosis. *Fix:* control glare at source — external shading, light shelves, glazing selection, screen orientation — and specify usable intermediate positions.

THE SEALED CEILING. Acoustic treatment that deletes the thermal mass strategy. *Damages* `C` (peak load, plant size, energy). *Signature:* the thermal model assumed exposed soffit and the reflected ceiling plan shows continuous plasterboard; nobody compared the two drawings. *Fix:* S16.11 — baffles and rafts, coupled percentage recorded. Its sibling is the **carpeted slab**, the same deletion from below: count only the coupled face, per S16.12.

THE 5% FLOOR IGNORED. An operator treating residual thermal complaint as a personnel problem. *Damages* `C`, morale, and eventually u if the complainants are a non-random group. *Signature:* "we've optimised the setpoint, some people are never happy." *Fix:* accept the PPD floor, stop optimising the single number, buy dispersion coverage per S16.8.

THE DRESS CODE AS THERMAL POLICY. A handbook clause disabling the free adjustment channel. *Damages* `C`. *Signature:* clo uniform across occupants and season; cooling load peaking in a month when nobody would voluntarily wear a jacket. *Fix:* relax the clothing constraint before spending on plant. **The cheapest thermal intervention in the building, and not an engineering decision.**

THE CO2 OVERCLAIM. A ventilation business case built on the cognitive-performance literature. *Damages* your credibility, and the project when someone checks the citations. *Signature:* a slide asserting a percentage decision-making decrement at a stated ppm, with no mention of the failed replications. *Fix:* §16.5 — re-argue on symptom prevalence, infection risk and perceived air quality. Its mirror image is **the unverified ventilation rate**: nobody can state the delivered L/s per person, no CO2 log exists, and the only evidence is a commissioning certificate predating the partition layout. *Fix:* a logger and two weeks, per S16.10.

THE COLD HANDRAIL. Bare metal where skin meets building in an unheated space. *Signature:* people use the stair without touching the rail — also a safety finding. *Fix:* S16.13; the calculation takes two minutes.

THE LAMINATE DECADE. Layered materials specified on first cost. *Damages* `delta` (appearance), then `eta` (the place stops reading as cared-for), then `C` (step-function replacement, always deferred). *Signature:* year-1 photographs look better than year-8 on identical maintenance. *Fix:* S16.14. **The instrument that detects this is a camera on a tripod, once a year.**

Worked Example

The system

A four-person shared workroom in a research institute. 42 m² (6.0 x 7.0 m), 2.9 m to soffit, V = 121.8 m³, internal surface 159.4 m². One external wall, 6.0 m wide, north-facing, with 8 m² of air-filled double glazing (U = 2.8 W/m²K), head height 2.4 m. Exposed concrete soffit, vinyl on a concrete slab, plastered blockwork. One wall thermostat serving this room and two others. Six distinct specialisms across the occupants and their regular visitors.

Step 1 — Parameters before

Surveyed per Chapter 24, Paper E §5.1 choices declared:

- **`Q`** — normalised count of live shared methods, tools and working agreements in use. Measured `Q_0 = 4.0`.
- **`N = 6`**, non-redundant by Chapter 7's collision test. `ln N = 1.7918`.
- **`U`** — disagreement audit, `u = 0.80`, so `N_eff = 5.00` and `S = ln 5 = 1.6094`.
- **`eta = 0.28`**, durable structure per unit surplus over six quarters.
- **`delta = 0.070`** per month.
- **`C`** — **calibration declared:** `C = phi . eta . ln N`, `phi` a dimensionless fraction of the ceiling; tallied `phi = 0.62`.

```
eta*lnN = 0.5017  <- ceiling     eta*S = 0.4506  <- input     delta*Q = 0.2800
C = 0.62 x 0.5017 = 0.3110
h = 0.5017 - 0.2800 - 0.3110 = -0.0894
```

`h < 0`. In this model the room's order decreases regardless of how its occupants are arranged (Paper E, Corollary 3.4). You may not build here; you must cut `C`, raise `eta`, lower `delta`, or shed `Q`. Steady state `Q_inf = (0.4506 - 0.3110)/0.070 = 1.994`, against a current 4.0: the room is on a trajectory to lose half its stored order.

Step 2 — The substrate survey

Light. Desk illuminance 480 lux — compliant. Vertical at seated eye height, facing the work direction: 130 lux, mEDI 95 at midday. **Against a 250 mEDI target the room delivers 38%.** Average daylight factor 1.4%; useful daylit depth at a 2.4 m head is 4.8–6.0 m against a room depth of 7.0 m. North aspect: no direct sun, no glare, no blind use — the one thing this room got right, by accident of orientation.

Thermal. Air temperature at the thermostat 21.0degC. Globe measurement at the window desk on a 0degC day: operative **18.9degC**, a 2.1 K deficit, consistent with the computed 13.4degC inner glazing surface. Two of four desks are within 1.5 m of the glazing; personal channels available: **zero**. The shared thermostat logged **3.8 contested adjustments per day** across its three rooms, at a mean of 11 minutes of consequent discussion and counter-adjustment.

Air. CO2 logged three weeks; afternoon plateau with four occupants **1,290 ppm**. Inverting, `Q = 0.005 x 4 / (870 x 10⁻⁶) = 23.0 L/s` total, 5.7 L/s per person. **Ventilated almost exactly to standard and sitting at 1,290 ppm — §16.5's tension in a real measurement.** RH ranged 22–38% across the heating season, below the 30% floor for much of it.

Acoustics. Measured RT60 at 500 Hz: **2.6 s** — between the computed all-hard 4.41 s and a treated room, because furniture and occupants absorb. Four people cannot hold two conversations here; they hold one, and the quietest of them contributes least.

Material. Laminate desktops in year 6 of an 8-year life, edges lifting on two; powder-coated steel shelving; vinyl floor. Every visible surface is layered: **improving-with-wear fraction 0%.** Contact temperature of the steel sill, at 18degC surface: 20.1degC.

Step 3 — Diagnosis

The room fails five of the six thermal variables, the vertical light plane, humidity, reverberation and material class; it passes horizontal illuminance, glare and ventilation rate.

The binding constraint is `C`, and the substrate generates it: 2.1 K of radiant deficit at half the desks driving a contested shared control; 2.6 s of RT60 suppressing three of four voices; a step-function replacement liability two years out. `eta` is also low at 0.28, and §16.7's mapping says why — reverberation suppresses the turn-taking that produces durable agreement, and the window desks with the best light and view are avoided in winter.

Step 4 — Intervention

1. **Acoustic, mass-preserving (S16.11).** 17 m² of 50 mm mineral wool baffles hung clear of the soffit, a 10 m² wall panel, carpet tile replacing vinyl. Computed `A = 38.64 m² Sab`, **RT60 = 0.51 s**, with **60% of the soffit still thermally coupled**.
2. **Glazing (S16.7).** Double low-E argon (U = 1.4) replaces the air-filled unit. Inner surface at 0degC outdoors rises from 13.4degC to **17.2degC**; local operative deficit at the window desks falls from 2.1 K to under 1.0 K, inside the S16.7 limit.
3. **Personal channels (S16.8).** Four desk fans and two low-wattage radiant foot panels — two channels per position, personal in effect.
4. **Shared control (S16.9).** The room gets its own zone valve, so the contested control serves four people instead of twelve. Logging retained.

5. **Vertical light (S16.3, S16.4).** Two indirect uplights and a wall-wash at 4000 K, dropping output and colour temperature after 17:00, targeting 250 mEDI at eye height in daytime hours.
6. **Humidity.** Trim humidification on the serving air-handling unit to hold RH >= 30%.
7. **Material and monitoring (S16.14).** As the laminate tops expire, replace with oiled solid timber; timber sleeves on the steel sill raise contact temperature at 18degC from 20.1degC to about 28degC. Ten taped camera positions, photographed annually.

None of this changes the roster, the specialisms, or the work.

Step 5 — Parameters after

Re-measured at four months:

```
eta:   0.28 -> 0.44   turn-taking recovered; window desks occupied year-round
delta: 0.070 -> 0.055 solid timber and carpet tile decay slower than lifting laminate and
                scuffed vinyl; measured from the photo series plus replacement records
phi:  0.62 -> 0.38   contested adjustments 3.8/day -> 0.4/day; deferred liability removed
N, u:  unchanged at 6 and 0.80, so S = 1.6094 unchanged

eta*lnN = 0.7884    eta*S = 0.7082    delta*Q = 0.2200    C = 0.38 x 0.7884 = 0.2996
h = 0.7884 - 0.2200 - 0.2996 = +0.2688
Q_inf = (0.7082 - 0.2996) / 0.055 = 7.429        (was 1.994)
```

Headroom moves from -0.089 to +0.269, from "no arrangement of people prevents decline" to a positive maximum one-step increment. **A 3.7-fold increase in the sustainable level of stored order, from baffles, glazing, four fans, two panels, a zone valve, two uplights and a change of worktop material.**

Step 6 — The transient

`tau = 1/delta = 18.2 months`; half-life `ln 2 / (-ln(1-0.055)) = 12.3 months`; 95% settling at `3 tau ~ 55 months`.

Do not evaluate this room on `Q` at four months — it has moved perhaps a fifth of the way. Evaluate the *inputs* at four months (`C`, RT60, mEDI, operative temperature at the seat, contested-adjustment count), all of which respond immediately because they are the drive, not the state. Evaluate `Q` at three years. This is Chapter 2's step response in a workroom.

Step 7 — What would falsify this

1. **`phi` does not fall.** Contested adjustments above 2/day after zoning and personal channels means the dispute was never thermal. Diagnostic: log which occupant initiates, and whether their preferred temperature actually differs from delivered.
2. **`eta` does not rise.** Still near 0.28 at eight months with RT60 at 0.51 s means reverberation was not what suppressed conversion, and Chapter 4 is the next chapter, not this one.
3. **Postulate 16.2a fails.** Run Protocol 16.1. If distance from delivered operative temperature predicts neither hours present, nor optional attendance, nor retention, **the `N_eff` claim in §16.4 is false in this building and must be withdrawn from the case.** Report that result.

4. **The shuffled-input control (Chapter 25).** A matched room, same institute and period, given redecoration of equal cost and disruption but no change to RT60, glazing U-value, personal control count or vertical mEDI. If its Q rises comparably, the claimed mechanism is not the one operating.

Step 8 — The honest ledger

Shown: a substrate intervention changing no personnel and no category structure moved `eta` up 57%, `delta` down 21% and `C` down 4%, taking the room from negative headroom to positive. Not shown: that these are your numbers. `eta`, `delta`, `u` and `phi` were measured for one room; the calibration was declared, not derived; the σ of thermal preference in §16.4 was assumed and flagged; and the `eta` gain is attributed to acoustics and light on the basis of §16.7's mapping, a modelling assumption that falsification item 4 exists to break.

Two things stayed broken, recorded rather than smoothed over:

- **S16.1 not met.** The room is 7.0 m deep against a useful daylit depth of 4.8–6.0 m at a 2.4 m head. The back strip is electrically lit all day, and no lighting scheme changes that. Remedy requires a second aspect or a shallower room.
- **S16.10's tension unresolved.** The room sits at 1,290 ppm while meeting the ventilation standard exactly. Ventilation was not increased, because §16.5 says the cognitive case is not sound and the symptom and infection case was not made for this room. **The project record reads: ventilation compliant, CO2 1,290 ppm, no action, reason stated.** That is what an honest specification looks like when the evidence does not support the intervention you would have preferred to sell.

Chapter 17 . The Table With No Head

Governance is usually written down. This chapter is about the part that is built instead.

A table with a head is a governing instrument. It assigns one seat the sightlines, the acoustic advantage and the conventional right to open and close, and it assigns every other seat the complementary role. **The function of everyone below the salt becomes agreement — and agreement between aligned instruments is echo, not evidence.** That is not a metaphor. Chapter 7 priced it: when the members' contributions become parallel, u goes to zero, `N_eff` goes to one, and the surplus of the whole assembly is `ln 1 = 0`.

The chapter's claim is that this is fixable in joinery. Three arguments:

1. **The damage from deference is quartic, not linear.** Mild deference is nearly free. Strong deference is catastrophic. §17.2 computes the curve.
2. **Geometric equality has an exact definition you can compute from a plan**: the variance of visual centrality across seats. It is zero for rotationally symmetric arrangements and large for a long boardroom table. §17.3.
3. **Three independent constraints — proxemic, acoustic, and coordination cost — converge on the same table size.** None of them knows about the other two. §17.4.

The Law

What is proved

From Paper A, Theorem 2.1 and the effective-count identity:

`u(x,y) = 1 − (x.y)^2` `N_eff = g(u) = 1 + (N−1) u` `S = F(u) = ln N_eff`

with `0 <= F(u) <= ln N` (Paper D, Corollary 2.2), `F = 0` if and only if the vectors are parallel, and `F = ln N` if and only if they are orthogonal (Paper A, 4.3). And from Paper E §3 and §6.4:

`Q_{t+1} = (1−delta) Q_t + eta S_t − C_t` `h_t = eta ln N − delta Q_t − C_t`

with steady state `Q_inf = (eta F(u*) − c*) / delta` (Theorem 3.5), which **the theorem allows to be negative, and does not hide when it is.**

Two further proved results are load-bearing here. **Concavity (Paper A, 3.3):** the second genuinely different member is the largest single gain the assembly will ever buy, and the fortieth is nearly free of value. **The logarithmic ceiling:** doubling the number of distinct categories adds `ln 2 ~ 0.693` no matter where you start, so you cannot grow a committee out of a deference problem.

Two results from CANON §3 are the empirical anchors. **Condorcet's jury theorem requires independence**: collective accuracy exceeds individual accuracy only when errors are uncorrelated, and a chair who aligns the room has correlated them. **Woolley and colleagues (*Science*, 2010)** found a collective intelligence factor c predicting group performance across varied tasks, not predicted by members' average or maximum IQ, and correlating with social sensitivity, **equality of conversational turn-taking**, and the proportion of women. Turn-taking equality is the driver this chapter can build for.

What is postulated

POSTULATE 17.1 (Deference Geometry). Let h be the chair's position vector and let each other member have a private view `w_i`, unit, mutually orthogonal, and orthogonal to h. Model the contribution actually made by member `i` under deference as

`v_i = a h + sqrt(1 − a^2) w_i  a in [0,1]  a = the alignment coefficient`

Then `v_i . v_j = a^2` for `i != j`, so `u_ij = 1 − a^4`, and the assembly's effective count is `N_eff = 1 + (n−1)(1 − a^4)`.

This is a modelling postulate and it is badged. What is postulated is the *form* of deference — that it acts as a rotation of each contribution toward the chair's, by a common coefficient. What follows from that form is exact arithmetic. The postulate is wrong in detail: deference is not uniform across members, private views are not exactly orthogonal, and a may vary by topic. It is nonetheless the simplest model with the right qualitative behaviour, and §17.2 gives the measurement that would estimate a in a real room.

POSTULATE 17.2 (Visual Centrality). The share of a member's contributions that reach the assembly as contributions — rather than as asides — increases with the number of other members within a comfortable head turn of their forward gaze. **Badged.** It is consistent with the turn-taking literature and with direct observation, and §17.3 defines the quantity so that it can be computed from a layout and correlated against measured speaking share.

POSTULATE 17.3 (Continuous Enforcement). A constraint embodied in furniture is enforced on every occasion of use, at zero marginal administrative cost, without an enforcer. **Badged**, and §17.6 states its limit honestly: geometry raises the cost of dominance, it does not forbid it.

Honesty clauses carried into this chapter

- Paper A is pure mathematics and makes **no empirical or metaphysical claim**. The deference arithmetic of §17.2 is a consequence of Postulate 17.1, not a finding about committees.
- As deployed, the surplus functional is a **lower bound** on the ensemble's information.
- **The triangle inequality fails for geometric triples (Paper A §5).** Subadditivity holds on the parameter u only. Do not chain "distances" between members' positions.
- **Paper D declines to identify surplus with Shannon entropy or mutual information in general.** §17.5 uses the exponential of a speaking-share entropy as a diagnostic. That is a Hill number, a standard diversity index, borrowed here by formal analogy. **It measures turn-taking equality. It does not compute surplus, and the resemblance in form is not a theorem.**
- Woolley's c was measured in short-task laboratory settings. Extrapolating it to a standing council is a modelling assumption, not a result.
- **The acoustic route to c is NOT CONFIRMED, and this chapter carries that status.** Chapter 14 Postulate 14.2 chains acoustic treatment to c through Woolley's turn-taking correlation, and §14.13 states the falsification test that the chain requires. **The worked case in Chapter 14 ran that test and reported outcome (b) or a near miss of it: the mediator moved decisively (`N_eff_turn/n` 0.604 -> 0.812, matched control room flat at 0.638 -> 0.650), the outcome moved only slightly, and the causal claim was not confirmed.** Wherever this chapter buys turn-taking equality with geometry and acoustics (§17.5, S17.7, S17.8), it is buying the *mediator*. **Do not report a geometry or acoustics change as a c-factor intervention.** Report what is established: reduced listening effort, recovered access

to the `ln N` ceiling for members the room was excluding, and a measured rise in `N_speak/n`. Those are worth having on their own account and require no causal claim.

- **No protocol in this chapter may remove a person's ability to decline to speak** (CANON §5.4). Every speaking-order protocol here includes an unconditional pass.

The Engineering

17.1 Below the salt

The phrase is literal. On a medieval English hall table the salt cellar stood as a marker; the host and honoured guests sat above it, on the dais end, and the household and lesser guests sat below. Seat position encoded standing, and everyone in the room could read it without being told.

A head of table does four things mechanically, and each is a geometric fact rather than a social convention:

1. **It is the only seat with every other seat inside its forward cone.** §17.3 computes this: at an eighteen-place boardroom table the head sees fifteen of the other seventeen within a comfortable head turn, while a mid-side seat sees five.
2. **It has the best acoustic position.** It is the seat everyone else turns toward, so its speech arrives on-axis at every listener, and the listeners' speech arrives on-axis at it.
3. **It holds the conventional right to open, allocate and close.** Speaking first sets the anchor; speaking last has the final word; controlling allocation controls who is heard.
4. **It is the default addressee.** Contributions are made *to* the head rather than to the room, which converts a deliberation into a series of bilateral reports.

Items 1 and 2 are built. Items 3 and 4 are convention, but convention that the geometry continuously reinforces, because it is physically easier to address the person you are facing.

The consequence is Chapter 11's parallel wall assembled out of chairs. Once the room's job is to respond to one position, the contributions rotate toward that position, and the arithmetic of Chapter 7 applies without amendment: **ten aligned members are one member, and the assembly's report is not eleven independent judgements but one judgement repeated eleven times.** The committee then reports unanimity, and unanimity is read as confirmation, and it is not confirmation because the instruments were not independent.

17.2 The arithmetic of deference

Under Postulate 17.1, with n members (one chair and n−1 others) and pairwise $u = 1 - a^4$:

a	a^4	u	N_eff (n = 12)	S = ln N_eff	% of ln 12
0.00	0.0000	1.0000	12.000	2.4849	100.0
0.30	0.0081	0.9919	11.911	2.4775	99.7
0.50	0.0625	0.9375	11.312	2.4259	**97.6**
0.60	0.1296	0.8704	10.574	2.3584	94.9
0.70	0.2401	0.7599	9.359	2.2363	90.0
0.80	0.4096	0.5904	7.494	2.0142	81.1
0.90	0.6561	0.3439	4.783	1.5650	**63.0**

a	a^4	u	N_eff (n = 12)	S = ln N_eff	% of ln 12
0.95	0.8145	0.1855	3.040	1.1120	44.8
0.98	0.9224	0.0776	1.854	0.6173	24.8
1.00	1.0000	0.0000	1.000	0.0000	**0.0**

Read the shape rather than the rows. **At `a = 0.5` — members who are half-aligned with the chair — the assembly keeps 97.6% of its surplus. At `a = 0.9` it keeps 63%. At `a = 0.95` it keeps 45%.** The loss is negligible until it is severe, because `u = 1 − a⁴` and the fourth power does nothing until a approaches one, then does everything.

That is the chapter's first practical finding and it cuts both ways:

- **Ordinary politeness is nearly free.** A room where people are mildly influenced by the chair loses a few per cent. Do not spend governance effort on it.
- **A room where members are 90% aligned has lost more than a third of its assembly**, and a room at 95% has lost more than half — while still tabling twelve names, twelve salaries, and twelve signatures on the minute. **Nothing in the paperwork records the loss.**

The same computation at other sizes:

a	n = 6 (% of ln 6)	n = 12	n = 25 (% of ln 25)
0.50	97.0%	97.6%	98.1%
0.70	87.5%	90.0%	91.9%
0.90	55.8%	63.0%	69.1%
0.95	36.6%	44.8%	52.7%
1.00	0.0%	0.0%	0.0%

A larger assembly is slightly more robust in *percentage* terms, and this is the trap: adding members appears to help and barely does, because the ceiling it is protecting is `ln n`, which grows logarithmically. **Twenty-five deferential members at `a = 0.9` deliver 2.225 nats against twelve independent members' 2.485. Twice the room, twice the cost, less surplus.**

The storage-law consequence Take a standing council: `n = 13`, `eta = 0.50`, `delta = 0.08` per month, `Q = 9.0`, and cost expressed by the calibration `C = phi . eta . ln n` with `phi = 0.45` (**calibration choice, declared**, per Paper E §5.1).

```
ln n = ln 13 = 2.5649        eta ln n = 1.2825        delta Q = 0.7200
C = 0.45 x 1.2825 = 0.5771
```

Condition	a	u	N_eff	S	eta.S	h	Q_inf
Head of table	0.90	0.3439	5.127	1.6345	0.8172	**-0.015**	3.00
No head	0.30	0.9919	12.903	2.5574	1.2787	**+0.114**	10.37

The deferential council sits at `h ~ 0`: maintenance only, no growth available, whatever it resolves to do. Removing the head takes it to `h = +0.114` and raises the sustainable level of

stored order from 3.0 to 10.4 — a factor of 3.5, with the same thirteen people, the same room, the same agenda.

The break-even alignment is worth computing, because it is a hard number a chair can be shown:

```
eta S = C   ->   S_crit = C / eta = 0.5771 / 0.50 = 1.1542
N_eff,crit = e^1.1542 = 3.1716
u_crit = (3.1716 - 1) / 12 = 0.1810
a_crit = (1 - 0.1810)^(1/4) = 0.9513
```

Above `a = 0.951`, this council destroys stored order every period it meets. It is not merely failing to add; the meetings are net-negative, because the coordination cost exceeds what a one-effective-member assembly can produce. That is Corollary 3.4 arriving in a committee room.

Measuring a a is not directly observable, and a chapter that gave a table without a measurement would be selling arithmetic. Two instruments:

Protocol 17.1 — the written-then-spoken divergence. Before any discussion, each member writes a position on a scaled question (a numeric estimate, a rating, a rank order). Collect. Hold the discussion. Collect a second written position. Compute the variance of the first set, `V_pre`, and of the second, `V_post`. Under Postulate 17.1 the alignment coefficient relates to the variance ratio as

```
V_post / V_pre  ~=  1 - a^2        so       a_max = sqrt(1 - V_post/V_pre)
```

Badge: that relation follows from the postulate's linear form, not from data. Both variances are directly measurable and the protocol costs five minutes. Its limitation is that legitimate persuasion by evidence *also* reduces post-discussion variance, so the expression above is an **upper bound** on a. Separate the two by asking, in the second collection, what changed the member's mind: a specific argument, or nothing they can name. Let `f` be the fraction of changed positions that no member could attribute to an argument. Then

```
a_min = sqrt( f * (1 - V_post/V_pre) )
```

Convergence nobody can attribute to an argument is the part that belongs to a. Report the interval `[a_min, a_max]` and carry its midpoint as the working estimate. **A point estimate of a reported without its interval is the most misleading number this chapter can produce.**

Protocol 17.2 — the addressee count. During the meeting, log for each contribution whether the speaker's gaze and body orientation are directed at the chair or at the table. A room in which over half of contributions are addressed to the chair is a room operating at high a, and the count requires one observer and a tally sheet.

17.3 Sightline reciprocity

The head-turn geometry of a circle Seat every member on a circle of radius `r`, facing the centre. For two members separated by angle φ around the circle, the head turn ψ that member `i` must make to face member `j` satisfies

```
cos psi = |sin(phi/2)|        equivalently        psi = 90 deg - phi/2
```

Verified numerically:

Separation ϕ	27.7deg	45deg	60deg	90deg	120deg	166deg	180deg
Required head turn ψ	76.2deg	67.5deg	60.0deg	45.0deg	30.0deg	6.9deg	**0.0deg**

Comfortable sustained head rotation is about ±45deg, with ±60deg reachable and beyond that requiring the torso.

The result is counter-intuitive and it matters. **At a round table the person directly opposite you is the easiest to face — zero head turn — and your immediate neighbour is the hardest.** At thirteen seats the neighbour sits 27.7deg away and needs a 76deg turn. For n = 13:

Seat offset	±1	±2	±3	±4	±5	±6
Separation	27.7deg	55.4deg	83.1deg	110.8deg	138.5deg	166.2deg
Head turn ψ	76.2deg	62.3deg	48.5deg	34.6deg	20.8deg	6.9deg
Class	torso	torso	workable	easy	easy	easy

A round table is a plenary instrument, not a conversational one. It is excellent at making each member addressable by the assembly and poor at supporting side conversation — which is exactly the trade a deliberating body wants, and exactly the wrong trade for a dinner. Design for the function.

Visual centrality, and the exact definition of geometric equality Define, for each seat `i`:

```
V_i = the number of other seats within a 45 degree head turn of i's forward gaze
```

`V_i` is computable from a furniture layout with a protractor, before anyone sits down. Now the definition this chapter is for:

Geometric equality is the condition Var(V_i) = 0.

Computed for real layouts, with seats 0.6 m out from the table edge and facing the table:

Layout	n	V per seat	mean	**variance**
Round, 12 seats, D = 2.67 m	12	7 everywhere	7.00	**0.00**
Regular 13-gon ring, D = 2.90 m	13	6 everywhere	6.00	**0.00**
Hollow square 2.4 x 2.4 m, 3 per side	12	7 everywhere	7.00	0.00
Rectangle 3.0 x 1.2 m, 5+5 sides + 2 heads	12	6,6,5,5,5,5,5,5,6,6,9,9	6.00	**2.00**

Layout	n	V per seat	mean	**variance**
Boardroom 6.0 x 1.4 m, 8+8 + 2 heads	18	5,5,5,5,6,6,7,7,7,7,6,6,7,8,5,5,5,5,	**15,15**	**9.06**

Read the last row. **At an eighteen-place boardroom table the two head seats each see fifteen of the other seventeen within a comfortable turn; the seats at the far ends of the long sides see five.** The head has three times the visual reach of the average member. That is not a rule of order, it is a fact about the shape of a rectangle, and it operates in every meeting held at that table, forever, without anyone deciding it should.

Compare the round table: variance exactly zero, by rotational symmetry. Every seat has the same view of the assembly. **There is no seat to take, so there is nothing to arrive early for.**

One honest qualification, and it is the kind that separates a computed claim from a slogan. The hollow square also scores zero variance at the ±45deg threshold — but that is a coincidence of the threshold, not a symmetry. Recompute at other thresholds:

Threshold	Round, n=12	Hollow square, n=12
±30deg	variance 0.00	variance **0.89**
±45deg	0.00	0.00
±60deg	0.00	**0.22**
±70deg	0.00	0.00

Only a rotationally symmetric arrangement has zero centrality variance at every threshold. The hollow square is a good second choice and it is not equal; its corner-adjacent seats and its mid-side seats are in genuinely different positions. If you specify a hollow square, say so, and record the variance at all four thresholds rather than the one that flatters it.

Protocol 17.3 — the centrality audit. Take the furniture plan. Mark each seat position and its facing direction. For each ordered pair, measure the angle between the seat's facing direction and the bearing to the other seat. Count those within 45deg. Report the vector `V`, its mean, and its variance at 30deg, 45deg and 60deg. **A layout whose variance you have not computed is a layout whose governance you have not specified.**

17.4 Sizing the table

Three independent constraints. They were derived for different reasons by people who did not consult each other.

Constraint 1 — proxemic Seat pitch along the table edge: 0.60 m is the minimum place setting, 0.70 m comfortable, 0.75 m generous. For a circular or regular polygonal table, `D = n s / π`, and the across-table distance between opposite members is `D`.

Hall's bands: personal 0.45–1.20 m, social 1.20–3.60 m, public beyond 3.60 m. **Across-table distance must stay inside the social band, or the far member is being addressed rather than conversed with.**

n	6	8	10	12	**13**	14	16	18	20	25
D at s = 0.70 m	1.34	1.78	2.23	2.67	**2.90**	3.12	3.57	4.01	4.46	5.57
Hall band	social	social	social	social	social	social	social	**public**	public	public

```
D <= 3.60 m (edge of social)          ->  n <= pi(3.60)/0.70 = 16.2
D <= 2.90 m (comfortable, facial detail) ->  n <= pi(2.90)/0.70 = 13.0
```

Thirteen seats at 0.70 m pitch give a 2.90 m table — exactly the comfortable social distance. The thirteen-sided plan is not a mystical number here; it is what falls out of anthropometry and Hall's bands, and it is stated as a derivation rather than a symbol.

A 2.90 m table has a great deal of unreachable centre. Comfortable reach from a seated position is 0.45–0.60 m, so make it a ring: at 0.60 m working depth the inner void is `2.90 − 2(0.60) = 1.70 m`, and the table's own footprint is 4.33 m². The ring is not a stylistic choice; a solid 2.9 m disc is 6.6 m² of which 2.3 m² can never be reached.

A measured artifact: the Winchester Round Table The great oak table hanging in the Great Hall at Winchester is about 18 feet in diameter and carries twenty-five painted places — Arthur and twenty-four knights. Compute:

```
D = 18 ft = 5.486 m          circumference = pi D = 17.24 m
17.24 / 25 places = 0.689 m per place
```

A fourteenth-century object is dimensioned at 0.689 m per seat, which is the modern comfortable place setting to within a centimetre. Anthropometry has not changed and neither has the answer.

But run the second calculation. Across-table distance is 5.49 m, which is well into Hall's public band. **By its own geometry the Winchester table is a ceremonial object and not a deliberation table**, and the round form does not rescue it: at 5.5 m the person opposite you is being addressed, not consulted. The lesson is precise and worth keeping: **roundness buys equality; it does not buy conversation. Conversation is bought by diameter.**

Constraint 2 — acoustic Speech level falls with distance. In the direct field, `L(d) = L(1 m) − 20 log10(d)`. Normal conversational effort is about 60 dBA at 1 m; relaxed about 55; raised about 66. Good speech intelligibility wants a signal-to-noise ratio of about 15 dB, and 10 dB is the adequate floor.

Across-table distance d	1.5 m	2.0 m	2.9 m	3.6 m	4.5 m	5.5 m
Level from a normal voice	56.5	54.0	**50.8**	48.9	46.9	45.2 dBA

Across-table distance d	1.5 m	2.0 m	2.9 m	3.6 m	4.5 m	5.5 m
Background needed for SNR 15	41.5	39.0	**35.8**	33.9	31.9	30.2 dBA
Background needed for SNR 10	46.5	44.0	40.8	38.9	36.9	35.2 dBA

At the 13-seat, 2.90 m table, a normally-spoken sentence arrives at 50.8 dBA, so the room's background noise must be at or below about 36 dBA — an NC/RC 30 room — for the quietest member to be heard without raising their voice. Go to twenty seats and 4.46 m and the requirement becomes 32 dBA, an NC 25 room, which is a concert-hall-grade mechanical specification and is rarely achieved in an office building.

Table size and the HVAC noise limit are the same constraint. The person who chose a twenty-seat table has specified the air handling, and does not know it.

Honest qualification. This is a direct-field calculation. In a real room the reverberant field adds level as well, but late-arriving reverberant energy smears rather than helps, so the direct calculation is the right conservative bound for *clarity*. Use it together with Chapter 14's RT60 target; neither replaces the other.

Constraint 3 — coordination cost Brooks: communication paths grow as `n(n−1)/2` while surplus grows as `ln n`.

n	6	8	10	12	13	16	20	25	30
Links n(n-1)/2	15	28	45	66	78	120	190	300	435
ln n	1.792	2.079	2.303	2.485	2.565	2.773	2.996	3.219	3.401
Links per nat	8.4	13.5	19.5	26.6	30.4	43.3	63.4	93.2	127.9

Maximising `eta ln n − c . n(n−1)/2` over `n` gives `n(n − 0.5) = eta/c`, so `n* ~ sqrt(eta/c)`:

η/c	25	50	100	200	400
Optimal n*	5.3	7.3	10.3	14.4	20.3

The convergence

```
Proxemic (facial detail, s=0.70):   n <= 13
Proxemic (edge of social):          n <= 16
```

```
Acoustic (achievable NC 30 room):   n <= about 13–16
Coordination cost, eta/c = 50..200: n* = 7 to 14
```

Three constraints derived from anthropometry, from acoustics, and from Brooks's arithmetic converge on eight to fourteen seats. None of the three knows about the other two. That convergence is the strongest result in this chapter, and it is the reason to distrust any deliberating body sized above about sixteen. Above sixteen you are not running a deliberation; you are running a plenary that will be decided somewhere else, by fewer people, in a smaller room.

17.5 Protocols

Geometry sets the field. Protocol sets the play. Four, in the order they are used.

Silent-write-first **The mechanism is precise.** An information cascade (Bikhchandani, Hirshleifer and Welch, 1992) requires that each actor observe the actions of earlier actors before acting. Writing first makes every member's first commitment simultaneous, so there are no earlier actions to observe. **The cascade is not discouraged; it is made structurally impossible for the first round.**

Procedure:

1. Question stated in writing, with the answer format fixed (a number, a rating, a ranking, or a short position statement — not an open essay).
2. **Two to five minutes of silence, everybody writing.** No talking, no screens shared, no glances at neighbours' pages. The chair writes too.
3. Simultaneous reveal: all positions posted at once, on a wall or a shared surface, unattributed if the group is new or the topic is sensitive.
4. Record `V_pre` — the variance across the written positions — before any discussion.
5. Discuss.
6. Second written round; record `V_post`.

The measurement is free and it is the group's own diagnostic. Illustrative: thirteen members writing estimates `3, 7, 2, 8, 5, 6, 4, 9, 3, 7, 5, 6,` 8 have mean 5.615 and variance 4.391. If post-discussion standard deviation falls to 0.6, the variance ratio is 0.082 — **91.8% of the group's independent variance has been destroyed**. If the members can point to the argument that did it, that is deliberation working. If they cannot, that is a cascade, and Protocol 17.1 estimates a ~ `sqrt(1 - 0.082) = 0.958`, which §17.2 says is past the break-even point.

Round-robin One pass, every member speaks once, no interruption, fixed time each. Two rules make it work:

- **The order is randomised each session**, or rotated by a rule nobody controls. A fixed order hands the anchoring advantage permanently to whoever is first.
- **A pass is unconditional.** "Pass" is a complete contribution and needs no reason. This is CANON §5.4 in a meeting: no protocol may remove a person's ability to decline.

Last-speaks-first inversion Where rank is present and cannot be removed, invert the speaking order: lowest rank first, highest last. The senior member's position is then stated after everyone else's is on the record, so it cannot anchor them. The practice is long-established in aviation crew

resource management, in military councils and in clinical teaching rounds, and it is presented here as **established practice rather than as a cited experimental result** — the mechanism it defeats (anchoring on an authoritative first statement) is well documented; the specific field effect sizes are not something this manual has measured.

Combined with silent-write-first, the inversion is largely redundant for the first round and remains valuable for the discussion, which is where anchoring returns.

Measuring turn-taking equality Woolley's driver is equality of conversational turn-taking, so measure it. Let `p_i` be member `i`'s share of speaking time. Compute the Shannon entropy of the share distribution and its exponential:

```
H = − sum_i p_i ln p_i          N_speak = exp(H)          E = H / ln n
```

`N_speak` is the **effective number of speakers**. It is a Hill number — a standard diversity index — and it is being used here as a diagnostic. **BADGE, and it is the same badge as CANON's: Paper D explicitly declines to identify surplus with Shannon entropy or mutual information. `N_speak` measures how equally the floor was shared. It does not compute `N_eff` and it does not compute surplus. The resemblance of form is a resemblance, not a theorem.**

Computed:

Case	n	H	E = H/ln n	N_speak	Gini
8 equal	8	2.079	1.000	8.00	0.000
8, one at 35%	8	1.912	0.920	6.77	0.225
8, two at 30% each	8	1.806	0.868	6.08	0.350
8, one at 50%	8	1.666	0.801	5.29	0.375
8, one at 60% + two silent	8	1.317	0.633	**3.73**	0.575
12 equal	12	2.485	1.000	12.00	0.000
12, chair at 40%	12	2.112	0.850	8.26	0.317
12, chair at 40% + 3 silent	12	1.921	0.773	**6.83**	0.467

A twelve-person meeting in which the chair takes 40% of the floor and three members say nothing has 6.8 effective speakers. Half the room is decorative, and the minute will record twelve attendees.

Report both `N_speak` and the Gini coefficient. Entropy is forgiving of a single dominator and harsh on silence; Gini is the reverse. Two indices disagreeing is information.

Instrument. A phone stopwatch and a tally, or an automated speaker-diarisation tool if the group consents to recording. Consent is not optional and a group that declines should be measured by hand or not at all.

What the chair still does Removing the head of the table does not remove the chair's work. Timekeeping, agenda-setting, recording the decision, and enforcing the protocols are all real and all necessary. **None of them requires the geometric privileges.** The chair can hold every one of those functions from an ordinary seat, and should, because a chair with the same sightlines as everyone else is a chair whose authority rests on the role rather than on the furniture.

One caution on a popular device: an *assigned* devil's advocate is weaker than genuine dissent. Charlan Nemeth's research on minority influence found authentic dissent to stimulate more and better divergent thinking than role-played disagreement. **Assigning the role is better than nothing and it is not a substitute for recruiting someone who actually disagrees** — which is Chapter 9's problem, not this chapter's, and cannot be solved with furniture.

17.6 Why furniture rather than policy

A rule requires an enforcer. The enforcer needs authority, discretion, attention and willingness, and every act of enforcement is a transaction that costs both parties something. In the law's terms, **a rule generates a stream of `C` for as long as it exists**, and the stream falls hardest on whoever is most willing to speak up — who, as Chapter 15 noted, is not a random sample.

A shape enforces on every occasion of use. It has no discretion, grants no exceptions, forgets nothing, retires from nothing, and asks nobody to be the person who raised it. Its marginal administrative cost is zero.

Furniture is the cheapest governance instrument ever built, and it enforces continuously without anyone administering it.

The comparison in units:

Instrument	Setup cost	Marginal cost per use	Decays without maintenance?	Who bears enforcement?
A rule in the terms of reference	low	one enforcement act each time	yes, quickly	a person, repeatedly
A trained facilitator	high	a salary, every session	yes, on turnover	the facilitator
A round table	one-off capital	**zero**	only physically, over decades	**nobody**

This is Star and Ruhleder's point about infrastructure: it becomes visible only on breakdown. A round table is invisible governance. Nobody in the room experiences it as a constraint, nobody resents it, nobody has to be reminded of it, and nobody has to be the person who enforces it — which is exactly why it works and exactly why it is never in the governance document.

The honest limit, and it must be stated as plainly as the claim. Geometry constrains; it does not determine. A sufficiently determined chair will dominate a round table by talking more, and a disciplined group will deliberate well at a rectangle. **What the geometry changes is the cost of dominance and the cost of resistance.** At a table with a head, dominating is free and resisting is expensive. At a table without one, dominating requires visible effort — you have to take more floor time, in front of people who can all see you doing it — and resisting requires only speaking, from a seat with the same view as everyone else's. That is a shift in the gradient, not a prohibition, and Chapter 20's argument applies: **a method that works by removing resistance has no coercive failure mode. Its worst outcome is that someone declines to use it.**

The Specification

S17.1 — No head. The deliberation table shall be circular or a regular polygon with n equal sides, seating one member per side or per equal arc. No seat shall differ from any other in its relation to the table's geometry.

S17.2 — Centrality variance. Computed per Protocol 17.3, `Var(V_i)` shall be **zero at the ±30deg, ±45deg and ±60deg thresholds**. Report all three. A layout achieving zero at one threshold only shall be recorded as unequal.

S17.3 — Seat pitch. Not less than **0.60 m** of table edge per member; **0.70 m** is the design default; 0.75 m where writing surfaces or documents are in use.

S17.4 — Diameter and seat count. `D = n s / π`. **`D` shall not exceed 3.60 m** (Hall's social limit). Where facial detail is required for the whole assembly, `D` shall not exceed **2.90 m**, which at 0.70 m pitch caps the body at **thirteen seats**.

S17.5 — Ring form. Where `D` exceeds 1.80 m the table shall be a ring, with a working depth of **0.45–0.60 m** and an open centre. Record the inner void diameter.

S17.6 — Upper bound on the body. No deliberating body shall exceed **sixteen** members. Above sixteen the assembly is a plenary and shall be documented as one, with its actual deliberating sub-body identified by name.

S17.7 — Acoustic budget. Background noise shall be not more than `L_speech(D) − 15 dB`, where `L_speech(D) = 60 − 20 log10(D)` dBA for normal vocal effort. At D = 2.90 m this is **36 dBA** (NC/RC 30). **The table diameter and the mechanical noise limit shall be signed off together, by the same person, on the same sheet.**

S17.8 — Reverberation. Meet the Chapter 14 RT60 target for the room volume; the direct-field budget of S17.7 does not substitute for it.

S17.9 — Chair seating. The person chairing shall occupy an ordinary seat, allocated by the same rule as every other seat, and shall not have a reserved position.

S17.10 — Seat allocation. Seats shall be unreserved, or allocated by a rule no member controls (rotation, draw). **No nameplates in fixed positions.**

S17.11 — Silent-write-first. Every substantive decision item shall open with 2–5 minutes of individual written positions, collected before any discussion. `V_pre` and `V_post` shall be recorded in the minute.

S17.12 — Cascade test. Where `V_post / V_pre < 0.25` and members cannot name the argument that moved them, the item shall be recorded as a suspected cascade and revisited at a later session.

S17.13 — Speaking order. The first round shall be round-robin in an order randomised or rotated by a rule no member controls. Where rank is present, the order shall run lowest rank to highest. **A pass shall be unconditional and shall require no reason.**

S17.14 — Turn-taking measurement. Speaking-time shares shall be measured at not less than one session in four. Report `N_speak = exp(H)` and the Gini coefficient. **Where `N_speak` falls below `0.7 n` for two consecutive measured sessions, the protocol has failed and shall be reviewed.** Measurement requires the consent of those measured.

S17.15 — Alignment estimate. Estimate a by Protocol 17.1 not less than annually, and compute `a_crit = (1 − (e^{C/eta} − 1)/(n−1))^{1/4}` for the body's own parameters. **Where the estimate exceeds `a_crit`, the body is destroying stored order and shall be reconstituted or dissolved.**

S17.16 — Minority record. Any member may attach a dissenting position to the minute, of their own drafting, unedited. This shall require no permission and no second.

S17.17 — Declination. No provision of this specification shall be applied so as to compel a person to speak, to be recorded, or to be measured.

The Failure Modes

THE HEAD OF TABLE. A seat with all others in its forward cone and the conventional right to open and close. *Damages* u toward zero and therefore `N_eff` toward 1. *Signature:* the centrality audit shows non-zero variance; over half of contributions are addressed to one seat; `V_post/V_pre` below 0.25 with no attributable argument; the same person speaks first and last in every item. *Fix:* S17.1 and S17.9. The table is the intervention.

THE LONG BOARDROOM TABLE. The head-of-table failure in its most severe geometric form. *Signature:* centrality variance of 9 or more; head seats with three times the visual reach of mid-side seats; members at the far corners contributing least, consistently, regardless of who sits there. *Fix:* the shape. **When the same seats produce the least contribution across changes of occupant, the seat is speaking, not the person.**

THE OVERSIZED BODY. A deliberating group above sixteen. *Damages* `C` quadratically while `ln n` creeps; also u, because the real decision migrates to a smaller unrecorded group. *Signature:* `N_speak` far below n; a rising fraction of members who never speak; decisions arriving at the meeting already made. *Fix:* S17.6. Name the sub-body that is actually deciding and constitute it properly, or shrink.

THE INAUDIBLE FAR SEAT. A table wider than the room's noise floor supports. *Damages* u by Chapter 14's mechanism — the loudest voice wins by physics. *Signature:* members at maximum across-table distance speak least; requests to repeat cluster by distance; quiet-voiced members are systematically under-represented in the minute. *Fix:* S17.7 — reduce `D`, or reduce background noise, or both. Do not "encourage people to speak up": the Lombard effect raises the noise floor along with the voices.

THE RESERVED SEAT. A round table with nameplates. *Damages* u, by reintroducing rank into a geometry designed to remove it. *Signature:* the same members occupy the same arc every session; arrival order matters; new members ask where to sit. *Fix:* S17.10. **A round table with fixed places is a rectangle with extra steps.**

THE CASCADE. Independent private judgements collapsing during discussion without an argument that did it. *Damages* u within the session, and `delta`-like over sessions as the habit sets. *Signature:* `V_post/V_pre` below 0.25; unanimity arriving fast; members unable to state what changed their mind. *Fix:* S17.11 and S17.12. **Unanimity that arrives quickly is the diagnostic, not the achievement.**

THE DECORATIVE MINORITY. Members present, counted, and silent. *Damages* `N_eff` — they

are not in the effective count regardless of the attendance list. *Signature:* `N_speak` below `0.7 n`; a stable set of members with near-zero speaking share; the same names never appearing in the minute. *Fix:* S17.13's round-robin plus S17.11's written round, both of which put a contribution on the record from every member without requiring anyone to interrupt.

THE ASSIGNED ADVOCATE. A rotating devil's-advocate role standing in for real disagreement. *Damages* u, by giving the appearance of dissent while everyone knows it is performance. *Signature:* the advocate's objections are never adopted; the role is treated as a chore; disagreement outside the role is still absent. *Fix:* out of scope for furniture — recruit genuine u per Chapter 9. Keep the role, but stop counting it as diversity.

THE PROTOCOL THAT COMPELS. A speaking-order rule with no unconditional pass, or a measurement regime imposed without consent. *Damages* the manual's ethics before it damages any parameter. *Signature:* members reporting that they spoke because they had to. *Fix:* S17.13 and S17.17. This is a red line under CANON §5.4, not a tuning parameter.

THE UNMEASURED TABLE. A round table installed and never instrumented. *Damages* your ability to know whether any of this worked. *Signature:* nobody can state `N_speak`, `V_post/V_pre`, or the centrality variance. *Fix:* S17.2, S17.11 and S17.14. **A geometry claim you have not measured is a decorating decision with a theory attached.**

Worked Example

The system

A research institute's standing scientific council. Thirteen members drawn from nine specialisms plus four cross-cutting roles; meets monthly; sets research priorities and allocates internal funds. Currently meets around a 4.2 x 1.4 m rectangular table with a chair position at one end, in a 7.0 x 6.0 x 3.0 m room with a hard plastered finish and a ceiling-mounted fan-coil unit.

Operational choices declared per Paper E §5.1:

- **`Q`** — normalised count of live council-originated policies, standards and allocation rules still in force and actually used. Measured `Q_0 = 9.0`.
- **`N = 13`** distinct positions, verified non-redundant by Chapter 7's collision test. `ln 13 = 2.5649`.
- **`eta = 0.50`**, from attributable durable structure per unit surplus over three years.
- **`delta = 0.08`** per month.
- **`C`** — **calibration declared:** `C = phi . eta . ln n`.

Step 1 — Survey

Centrality (Protocol 17.3). Thirteen seats: six per long side, one at the chair's end, all seated 0.6 m out from the edge. Computed `V = [4, 4, 5, 5, 6, 6, 6, 6, 5, 5, 5, 5, 10]`, mean 5.538, **variance 2.095 at ±45deg**. **The chair sees ten of the other twelve within a comfortable head turn; the two seats at the far end of the long sides see four.**

Acoustics. Across-table distance for opposite members 2.6 m; the working long-diagonal path is 4.4 m. Measured RT60 at 500 Hz: **1.4 s**. Measured background noise with the fan-coil running:

42 dBA. A normal voice delivers `60 − 20 log10(4.4) = 47.1` dBA at that range, so SNR 15 requires a background of **32.1 dBA**. **The room is 10 dB over budget on its long path.**

Addressee count (Protocol 17.2). Over three sessions, **61%** of contributions were directed at the chair rather than the table.

Turn-taking. Measured with consent over three sessions: chair 34% of speaking time, two senior members 13% each, three members at or near zero. Computed `N_speak = 6.9` against `n = 13`, so `N_speak / n = 0.53` — **below the S17.14 floor of 0.7.**

Alignment (Protocol 17.1). Two scaled items per session, four sessions. Mean `V_post / V_pre = 0.14`. In structured follow-up, members could name a specific argument in 3 of the 8 items, so the unattributed fraction is `f = 5/8 = 0.625`.

```
a_max = sqrt(1 − 0.14)              = 0.927
a_min = sqrt(0.625 x (1 − 0.14))    = 0.733
```

Working estimate `a = 0.83`, interval [0.73, 0.93].

Step 2 — Parameters before

```
n = 13,  a = 0.83   ->   u = 1 − 0.83^4 = 0.5254
N_eff = 1 + 12(0.5254) = 7.305        S = ln 7.305 = 1.9886
eta ln n = 0.50 x 2.5649 = 1.2825    eta S = 0.9943    delta Q = 0.08 x 9.0 = 0.7200
phi tallied at 0.45   ->   C = 0.45 x 1.2825 = 0.5771

h     = 1.2825 − 0.7200 − 0.5771 = −0.0146
Q_inf = (0.9943 − 0.5771) / 0.08 = 5.21
```

At the interval's ends, `Q_inf` runs from **6.88** (a = 0.73) down to **1.49** (a = 0.93). Every one of those is below the current `Q` of 9.0.

h is negative. This council cannot grow stored order in any month, by any decision it takes, and it is losing what its predecessors built. The break-even alignment:

```
S_crit = C / eta = 1.1542      N_eff,crit = 3.1716      u_crit = 0.1810
a_crit = (1 − 0.1810)^0.25 = 0.9513
```

The working estimate `a = 0.83` sits below `a_crit = 0.951`, so the council is not yet net-destructive — but the interval's upper end, `a = 0.93`, is within 0.02 of it. **Report the interval. A point estimate here would be the most misleading number in the survey.**

Step 3 — Intervention

Nothing changes the membership, the terms of reference, or the agenda.

1. **A thirteen-sided ring table (S17.1, S17.4, S17.5).** `D = 13 x 0.70 / π = 2.90 m`; working depth 0.60 m; inner void 1.70 m; table footprint 4.33 m². Centrality variance **0.00 at ±30deg, ±45deg and ±60deg**; every member sees six of the other twelve within a 45deg turn.
2. **Room (S17.7, S17.8).** Table plus 1.2 m clear circulation all round needs a **5.30 m** clear diameter, so a 5.3 m square room minimum. Taking 5.3 x 5.3 x 3.0 m: V = 84.3 m³, S = 119.8 m². For RT60 = 0.5 s, required absorption `A = 0.161 x 84.3 / 0.5 = 27.1 m²`

Sab, a mean coefficient of 0.227 — achieved with an absorbent ceiling raft plus about 8 m² of wall panel, with the thirteen occupied seats contributing 5.85 m² Sab of it unaided. Fan-coil replaced with a low-noise unit and attenuators: measured background **34 dBA** against the S17.7 budget of `60 − 20 log10(2.90) − 15 = 35.8 dBA`. **Inside budget, where the old room was 10 dB over.**

3. **Chair seating (S17.9, S17.10).** No reserved position, no nameplates; seats drawn each session.
4. **Silent-write-first (S17.11) and cascade test (S17.12)** on every substantive item, with `V_pre` and `V_post` in the minute.
5. **Round-robin, randomised order, unconditional pass (S17.13)**, with lowest-to-highest rank ordering on items where rank is present.
6. **Turn-taking measured quarterly with consent (S17.14).**
7. **Minority record (S17.16)** available to every member without permission or second.

Capital cost: one table, one ceiling raft, wall panels, and a fan-coil replacement.

Step 4 — Parameters after

Re-measured over the following four sessions:

```
addressee count:  61% to the chair  ->  24%
V_post / V_pre:  0.14          -> 0.52, argument nameable in 7 of 8 items, so f = 0.125
alignment a:      [0.73, 0.93]      ->  a_max = sqrt(0.48) = 0.693
              mid 0.83           a_min = sqrt(0.125 x 0.48) = 0.245    mid 0.47
N_speak:         6.9  (0.53 n)     ->  10.8   (0.83 n, above the S17.14 floor)
phi:          0.45          -> 0.38  (fewer repeat items, fewer reopened decisions)
```

```
a = 0.47  ->  u = 1 − 0.47^4 = 0.9512
N_eff = 1 + 12(0.9512) = 12.414     S = ln 12.414 = 2.5189
eta S = 1.2594        C = 0.38 x 1.2825 = 0.4873       delta Q = 0.7200

h     = 1.2825 − 0.7200 − 0.4873 = +0.0751
Q_inf = (1.2594 − 0.4873) / 0.08 = 9.65
```

At the interval's ends, `Q_inf` runs from **9.92** (a = 0.245) to **8.44** (a = 0.693) — the whole interval now sits at or near the council's current holdings, where before the whole interval sat below them.

Headroom crosses from -0.015 to +0.075, and the sustainable stored order rises from 5.21 to 9.65. The council stops losing what it inherited.

Note what did *not* happen. `N` is still 13, `eta` still 0.50, `delta` still 0.08. **The entire gain came from `u` — 0.525 to 0.951 — plus a modest fall in `C`, and `u` moved because the furniture and the protocols stopped rotating everyone's contribution toward one seat.**

Step 5 — The transient

`tau = 1/delta = 12.5 months`; 95% settling at about `3 tau = 38 months`. From Theorem 3.1 with `Q_0 = 9.0` and `Q_inf = 9.65`, the stored order barely moves — it was *already* near the new steady state, having been on its way down toward 5.21.

That is the correct and easily-missed reading: the visible effect of this intervention is a decline that does not happen. A council measuring only `Q` will see a near-flat line and conclude nothing was achieved. What was achieved is the difference between that flat line and a decay toward 5.21, which over three years is a large number that appears in no record.

Measure the inputs: addressee count, `V_post/V_pre`, `N_speak`, and background noise, all of which moved immediately, all of which are cheap, and all of which are the drive rather than the state. This is Chapter 2's lesson in its least intuitive form.

Step 6 — What would falsify this

1. **`N_speak` does not rise.** Below 0.7 n at two consecutive measured sessions, the protocols have failed and the geometry did not carry them. Diagnostic: check whether the chair is still allocating the floor outside the round-robin.
2. **`a` does not fall.** If the whole interval `[a_min, a_max]` stays above 0.80 after four sessions, the deference is not geometric in origin — it is contractual, or it is about funding — and no table fixes it. **Say so and stop, rather than buying more furniture.**
3. **`u` rises but `Q` does not stabilise.** If `Q` continues to decline through month 24, either `eta` is lower than measured or `delta` is higher, and the binding constraint was never `u`. Go to Chapter 4 or Chapter 3 respectively.
4. **The shuffled-input control (Chapter 25).** A comparable committee in the same institute, over the same period, given a refurbished room of equal cost and disruption but a rectangular table with a chair position retained, and no protocol change. **If its `u` and `N_speak` rise comparably, the mechanism claimed here is not the mechanism operating, and the effect belongs to attention rather than to geometry.** This is the control that decides whether this chapter is engineering.

Step 7 — The honest ledger

Demonstrated, within the model: an intervention touching only the geometry and the speaking protocols moved `u` from 0.525 to 0.951, took headroom from -0.015 to +0.075, and raised the sustainable stored order from 5.21 to 9.65 — a factor of 1.85 — with no change of membership.

Not demonstrated, and each of these is a real limit:

- **`a` is a modelled quantity, not an observed one.** Postulate 17.1 supplies its form; Protocol 17.1 supplies an interval, not a point; and the interval's lower end rests on attributing unexplained convergence to deference, which is a judgement made by the people being measured. The intervals here are wide — [0.73, 0.93] before, [0.25, 0.69] after — and they are reported wide rather than narrowed by assumption.
- **The calibration `C = phi . eta ln n` was chosen and declared**, not derived.
- **`N_speak` is a diversity index, not a surplus.** Paper D declines the identification and so does this chapter.
- **Woolley's `c` was measured on short laboratory tasks.** Its transfer to a monthly standing council is an assumption.
- **Geometry shifts the cost of dominance; it does not forbid it.** If this council's chair chooses to take 34% of the floor at the round table too, the round table will have failed, and the measurement in Step 4 is what will tell you.

And one requirement was not met and is recorded rather than relaxed: **S17.7 is satisfied for the across-table path at 2.90 m and only just — 34 dBA measured against a 35.8 dBA budget, a 1.8 dB margin.** A single additional noise source in that room, a laptop fan or an unattenuated diffuser, puts it over. The project record reads: *S17.7 met, margin 1.8 dB, no further noise sources permitted in this room without re-measurement.* A margin that thin is a finding, not a pass.

Part IV — The Sociology

Chapter 18 . Collective Intelligence as an Engineerable Quantity

The Law

Three statements govern this chapter. They have different warrants, and the difference matters more than the content.

STATEMENT 1 — PROVEN (Paper A, Theorem 2.1; Paper E, §3). The surplus available from an interaction is

```
F(x,y) = f(u(x,y)),    u = 1 − (x . y)^2 = ||x ^ y||^2  in [0,1]
f(u)   = ln(1 + (N−1)u)
N_eff  = g(u) = 1 + (N−1)u
```

and the order it leaves behind accumulates as

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t ,    S_t = F(U_t) = ln(1 + (N−1)U_t)
```

with `0 <= S_t <= ln N` (Paper D, Corollary 2.2). This is mathematics. It makes no empirical claim whatsoever. Paper A is pure mathematics and asserts nothing about groups, people, or minds.

STATEMENT 2 — MEASURED (Woolley, Chabris, Pentland, Hashmi & Malone, *Science*, 2010). Across a battery of varied short group tasks — visual puzzles, brainstorming, collective moral reasoning, negotiation over limited resources, and a criterion task — the authors extracted **a single statistical factor, which they labelled *c*, that predicted group performance across those varied tasks.** Their central findings, stated exactly:

1. *c* is **not** predicted by the **average** intelligence of the group's members.
2. *c* is **not** predicted by the **maximum** individual intelligence in the group.
3. *c* **is** significantly correlated with the **average social sensitivity of members**, measured by the Reading the Mind in the Eyes test.
4. *c* **is** significantly correlated with the **equality in distribution of conversational turn-taking** — groups where a few people dominated the conversation scored lower.
5. *c* **is** significantly correlated with the **proportion of women in the group**, and the authors reported that this effect was largely **mediated** by social sensitivity rather than standing as an independent driver.

The study ran roughly 699 individuals in groups of two to five, across two studies. The reported correlations were statistically significant but modest in magnitude — normal for this literature, and a limit we return to.

STATEMENT 3 — POSTULATED, BADGED HERE IN THE OPEN. This manual treats the three Woolley correlates as *operating handles on parameters of the storage law*. Specifically we postulate:

```
POSTULATE 18.1 (the c-bridge). BADGED: this is a modelling assumption, not a theorem
                               and not an experimental result.

  turn-taking equality          acts primarily on   U_t   (realised non-
parallelism)
 social sensitivity       acts primarily on  eta  (conversion of interaction to structure)
 task/room/protocol design  acts on         C_t  and on which U_t is reachable at all
```

Postulate 18.1 is the whole engineering content of this chapter and it is not proved. It is chosen because it is coherent with both the mathematics and the measurement, falsifiable by the Chapter 25 protocol, and productive of buildable specifications. It is not entailed by Statement 1 or Statement 2. Say so out loud whenever you lean on it, exactly as this book does here.

HONEST LIMIT, stated at the top rather than buried. *c* was measured in **short-task laboratory settings** — groups assembled from strangers, working for a session, on tasks with scorable outcomes. **Extrapolating *c* to a standing institution operating over years is a modelling assumption, not a demonstrated result.** Nothing in Woolley et al. licenses the claim that a hospital department, a standards body, or a city agency has a stable *c*. What it licenses is the far weaker and still useful claim that *group performance has structure that is not the sum of individual ability, and some of that structure lies in observable interaction behaviour.* Everything beyond that is engineering judgement wearing its badge.

Carried forward from CANON §1: the surplus functional as deployed is a **lower bound** on the ensemble's information. We log the floor.

The Engineering

18.1 What the 2010 result actually converted

Before 2010 the sentence "that team is smart" was an aesthetic judgement. It had the grammar of a measurement and the content of a compliment. You could not act on it, because you could not say what would have to change for the sentence to become more true.

The Woolley result converted that sentence into an engineering brief. Not because it explained group intelligence — it did not — but because it identified **observable behavioural quantities that co-vary with group performance and that a designer can touch.** Average IQ is not something a designer can touch on Monday morning. The distribution of speaking turns in a ninety-minute meeting is. **A correlate you can build for is worth more to an engineer than a cause you cannot reach.**

The correlates split cleanly into two kinds, and the split is the beginning of the design work:

```
KIND 1 — COMPOSITIONAL     who is in the room
         social sensitivity (a member property, aggregated)
         proportion of women (a member property, aggregated)

KIND 2 — BEHAVIOURAL       what happens in the room
         equality of conversational turn-taking (an interaction property)
```

Kind 1 is a staffing lever: slow, expensive, ethically fraught, and in most institutions largely fixed on the timescale of any intervention you will actually run. Kind 2 is a protocol, furniture, and acoustics lever: fast, cheap, reversible, available this week.

The turn-taking finding is the one an engineer should care about most, not because it is the largest effect but because it is the only one that is a property of the arrangement rather than of the people. You can change the arrangement without changing anyone.

18.2 Turn-taking equality, made numerical

"Equality of conversational turn-taking" is not yet a number. Make it one.

Let a group of n members produce speaking shares `p_1, ..., p_n` over an observation window, with `sum p_i = 1`. Define the participation entropy and the normalised equality index:

```
H = − sum_i p_i ln p_i                    (nats)
E = exp(H)                                 effective number of speakers
J = H / ln n                               equality index, in [0,1]
```

`J = 1` means perfectly even participation. `J -> 0` means one voice.

`E = exp(H)` is an **effective count of speakers** — a headcount audited for the fact that a person who says nothing is not, for the purposes of the discussion, present.

Note the structural rhyme with `N_eff = e^{F}` from Paper A. **BADGE: the rhyme is structural, not an identity.** `N_eff` is computed from geometric non-parallelism between category vectors; `E = exp(H)` from participation shares. They are different quantities sharing the exponential-of-a-log-measure form. Do not conflate them. The reason to notice the rhyme is that it tells you what kind of quantity you are handling: both punish redundancy, and both say the same brutal thing — *a member who adds nothing distinct is not counted.*

Here are the numbers, computed, for a five-person group:

participation profile	shares	H (nats)	E = exp(H)	J = H / ln 5
perfectly equal	0.20 / 0.20 / 0.20 / 0.20 / 0.20	1.6094	5.0000	1.0000
mild skew	0.30 / 0.25 / 0.20 / 0.15 / 0.10	1.5445	4.6855	0.9596
one dominant voice	0.60 / 0.15 / 0.10 / 0.10 / 0.05	1.2014	3.3247	0.7465
two-way duel	0.45 / 0.45 / 0.04 / 0.03 / 0.03	1.0578	2.8800	0.6573
severe dominance	0.80 / 0.10 / 0.05 / 0.03 / 0.02	0.7420	2.1001	0.4610

(`ln 5 = 1.6094`.)

Read the third row. A single person taking 60% of the airtime in a five-person meeting reduces the effective number of speakers from five to **3.32**. You are paying five salaries and convening three and a third people. At 80% dominance you have **2.10** effective speakers — a conversation between one person and a rounding error, which everyone present knows without being able to say why.

Read the fourth row carefully. **The two-way duel scores worse than the single dominant voice**, at `E = 2.88` against `3.32`. Two loud people arguing is not twice as good as one loud person talking; it is worse, because the argument consumes the airtime the other three would need in order to enter, and it does so while producing the *appearance* of vigorous debate. The duel is the most flattering of the failure modes and one of the most expensive.

18.3 From equality to the storage law

Now the bridge, wearing its badge.

POSTULATE 18.2 (realised versus available non-parallelism). BADGED. A group has an *available* non-parallelism `u_pot` determined by who is on the roster — their training lineages, incentives, and priors. It *realises* only some of that in any given interaction. We model the realised value as

```
U_t = J_t * u_pot            J_t in [0,1] the turn-taking equality index
```

This is a modelling choice and not a derived result. It is chosen for its boundary behaviour: at `J = 1` the group realises everything on the roster; at `J -> 0` only one voice reaches the record, and a single voice has `u = 0` with itself, so `F = ln 1 = 0`. A monologue produces no surplus no matter whose monologue it is. The linear interpolation between those endpoints is the simplest defensible form and should be replaced the moment anyone measures a better one.

With `N = 5` categories and an available `u_pot = 0.60`:

profile	J	U = J u_pot	N_eff = 1 + 4U	S = ln(1 + 4U)	% of the ln 5 = 1.6094 ceiling
perfectly equal	1.0000	0.6000	3.4000	1.2238	76.0%
mild skew	0.9596	0.5758	3.3031	1.1949	74.2%
one dominant voice	0.7465	0.4479	2.7915	1.0266	63.8%
two-way duel	0.6573	0.3944	2.5774	0.9468	58.8%
severe dominance	0.4610	0.2766	2.1065	0.7450	46.3%

Between an even meeting and a dominated one, the same five people on the same roster deliver 1.2238 nats or 0.7450 nats of surplus — a 39% loss purchased entirely by who was allowed to talk. Nobody was hired or fired. Nobody's ability changed. The roster is identical. The arrangement moved and the output moved with it.

Notice what the table does *not* say. Even at perfect equality the group reaches only 76% of its ceiling, because `u_pot = 0.60` and not 1. **Equality is a spending efficiency, not a source.** If the five are graduates of the same programme with the same incentives, `u_pot` is near zero, `J = 1` multiplies near zero by one, and the perfectly equal meeting produces `ln 1 = 0`. Chapter 11 is about that case; this chapter assumes you have avoided it.

18.4 Social sensitivity as a conversion parameter

Turn-taking governs how much difference gets *emitted*. Social sensitivity governs how much of what is emitted gets *received, understood, and retained*. That is the definition of `eta`.

Recall Chapter 4: `eta` is the fraction of interaction surplus that becomes durable structure. A meeting with no artifact has `eta` near zero. But `eta` fails in a subtler way too, before any artifact

question arises: **a point that is made and not understood was never converted, however well it was minuted.**

Consider the mechanism concretely. A junior clinician says, hesitantly and in the wrong vocabulary, that a discharge criterion looks unsafe in one specific class of case. Three things can happen:

1. Nobody notices she spoke. `eta` contribution: zero. Her turn is in the transcript and out of the system.
2. Somebody notices, and misreads the hesitancy as low confidence in the *content* rather than low confidence in her *standing to speak*, and moves on. `eta` contribution: zero — and the group now has positive evidence that the criterion is fine, which is worse than no evidence.
3. Somebody reads the hesitancy correctly and says *"say more about which cases."* The objection reaches the record. `eta` contribution: positive.

The difference between (2) and (3) is a perceptual skill: reading a state that was not stated. That is what the Reading the Mind in the Eyes instrument attempts to index, and it is why an `eta`-mapping is the natural reading of that correlate.

POSTULATE 18.3. BADGED. Average social sensitivity acts primarily on `eta`, the conversion of emitted surplus into retained structure. Not proved; adopted because it is coherent and testable.

The engineering consequence is important and slightly counter-intuitive. `eta` is, per Chapter 4, usually the **cheapest** parameter to raise — and the cheapest way to raise it is almost never to try to make people more perceptive. It is to **build prostheses that make perception unnecessary**:

perceptual demand	prosthesis that removes it	parameter touched
notice that someone wants to speak	round-robin protocol; explicit go-around	`eta`, and `J`
notice hesitancy behind a weak claim	silent written estimate before discussion	`eta`, and `u`
notice a dissent that was withdrawn	recorded minority position, required field	`eta`
notice a point already made and lost	decision log with an open-questions column	`eta`, and `delta`
notice that the quiet person disagrees	anonymous first-round elicitation	`eta`, and `u`

A protocol that makes a signal explicit converts a compositional advantage into an architectural one. You cannot reliably hire for social sensitivity, you cannot ethically screen for it, and you certainly cannot install it. You can install a go-around. The go-around does not make anyone more perceptive; it makes perception less load-bearing. This is the same move as Chapter 14's acoustic argument and Chapter 17's furniture argument, and it is the house method: **push the requirement out of the person and into the arrangement.**

18.5 The proportion-of-women finding, stated exactly and used carefully

The finding is: *c* correlated positively with the proportion of women in the group, and the authors reported the effect was **largely mediated by social sensitivity** — women in their sample scored

higher on the Reading the Mind in the Eyes instrument on average, and it was the sensitivity that carried the association.

Three things follow and one thing does not.

Follows. The mediation is the operative claim. A mediated correlate is a *marker* of the mediating variable in that population, not an independent lever. If the mediator does the work, the design response is to act on the mediator — which, per §18.4, means installing protocols that make social sensitivity less necessary, not sorting people.

Follows. A group composed entirely of one type on any dimension is, by Chapter 9's argument, already suspicious on `u_pot` grounds independent of anything in Woolley. Homogeneity of training lineage, of incentive, and of life experience all push `u` toward zero, and `u = 0` gives `N_eff = 1` and `S = 0` regardless of how sensitive anyone is. **The diversity argument in this manual does not rest on the 2010 paper and does not need it.** It rests on Paper A Theorem 4.3 and on Hong & Page (*PNAS*, 2004): under their conditions, a randomly selected diverse group of problem-solvers outperformed a group of the individually highest-performing solvers, because the high performers' heuristics were correlated.

Follows. The finding is a result about a sample, in a laboratory, at one time. It is a description, not a prescription.

Does not follow. That any individual's contribution can be predicted from their category. The correlation is between a *group-level proportion* and a *group-level score*, and it is modest. Using it to make claims about a person is a category error of exactly the kind the effective-count arithmetic exists to prevent: **a category is a direction in the space, and a person is not their category's basis vector.** Chapter 22 makes this the ethical spine of the book. Here it is enough to say that the honest use of this correlate is as evidence for a protocol, never as evidence about a human being.

18.6 *c* is not intelligence, and the difference is the effective count

The most important negative result in Woolley et al. is (1) and (2): neither the average nor the maximum member IQ predicted *c*. Engineers reading quickly convert this into "individual ability doesn't matter," which is false and unhelpful. The correct reading is sharper.

Ability is a property of a vector's *length*. Difference is a property of the *angle between vectors*. The surplus law depends only on `u = 1 − (x . y)^2`, which is a function of the angle alone — the vectors are unit-normalised, and length has been divided out. **The law is blind to ability by construction, and Woolley's negative result is what that blindness looks like when somebody goes and measures it.**

So: five brilliant people whose brilliance points the same way have `u ~ 0`, `N_eff ~ 1`, `S ~ ln 1 = 0`. They produce, jointly, the output of one brilliant person, minus the cost of convening them — by Chapter 5's Brooks arithmetic, `5 x 4 / 2 = 10` communication paths of overhead purchased for zero marginal surplus. The group is a strict loss against the best individual working alone. Anyone who has sat on a committee of experts has felt this and has usually blamed the personalities.

This is Ashby's Law of Requisite Variety (1956) in another vocabulary: only variety absorbs variety, and five instruments with identical response curves distinguish exactly what one distinguishes. It is also Condorcet's jury theorem with the fine print restored — collective accuracy rises toward

certainty with group size **only under independence**, and in the limit of perfect correlation a jury of a thousand has the accuracy of one juror. **Condorcet, Ashby, Hong & Page, and Paper A's extremality theorem are four statements of one fact, and Woolley's negative result is that fact showing up in a laboratory.**

Here is the marginal value of the next genuinely distinct member, at `u = 0.55`, computed:

N	S = ln(1 + (N−1)(0.55))	gain over N−1
1	0.0000	—
2	0.4383	+0.4383
3	0.7419	+0.3037
4	0.9746	+0.2326
5	1.1632	+0.1886
6	1.3218	+0.1586
8	1.5790	+0.1204
10	1.7834	+0.0970

The second genuinely different member is the largest single purchase you will ever make. The tenth buys 22% of what the second bought, and by Brooks costs `10 x 9 / 2 = 45` communication paths against the second member's one. **Concavity in surplus against quadratic growth in cost is the reason small groups win, and it is arithmetic, not culture.**

18.7 *c* is not morale, and *c* is not cohesion

Morale is how people feel about the group. Cohesion is how strongly members want to remain in the group and how strongly they align with it. Neither is *c*, and the second is frequently its enemy.

COHESION IS OFTEN THE ENEMY OF *c*, BECAUSE COHESION DRIVES u TOWARD ZERO.

That is the load-bearing sentence of this chapter, and it needs its mechanism spelled out, because stated bare it sounds like contrarianism and it is not. It is a consequence of what cohesion *is*.

Cohesion is a process that selects, over time, for alignment. It does so through at least five mechanisms, all ordinary and none malicious:

1. **Selection at entry.** Cohesive groups recruit people who will fit, and "fit" is measured by agreement on what the group cares about. `u_pot` reduction at the source, usually described as good hiring.
2. **Selection at exit.** Members whose priors diverge find the group uncomfortable and leave. `u_pot` reduction by attrition, usually described as natural turnover.
3. **Convergence under exposure.** Members who stay share framings until their priors align. This is Bikhchandani, Hirshleifer & Welch (1992): once enough public signals accumulate, rational individuals stop weighting private information and copy. **A cascade does not require anyone to be foolish or cowed. Copying is locally rational and it stops aggregation dead.**
4. **Suppression at emission.** Where private difference survives, the cost of voicing it rises, because voicing it reads as disloyalty. A `J` reduction: realised `U = J u_pot` falls even though `u_pot` has not.

5. **Reinterpretation of the record.** An objection raised and lost is remembered as having been answered. An `eta` failure applied selectively to disconfirming content.

Mechanisms 1-3 destroy `u_pot`. Mechanism 4 destroys `J`. Mechanism 5 destroys `eta` for exactly the inputs that would have corrected the group. **Cohesion attacks all three of the chapter's parameters simultaneously, which is why it is not merely one risk among several but the characteristic degenerative disease of successful groups.**

And note the cruel timing: cohesion is *produced by success*. A group that has worked well together becomes cohesive as a result. Its u then decays. Its performance follows, on a lag long enough that nobody connects the two. Chapter 27 catalogues this as the **Cohesion Death Spiral**.

Now put a number on the decay. Model `u_pot` as decaying geometrically under cohesion pressure:

```
u_{t+1} = (1 - lambda) u_t      lambda = cohesion-driven convergence rate per period
t_half  = ln 2 / ( -ln(1 - lambda) )
```

Starting at `u_0 = 0.60`, with `N = 5`:

lambda	half-life (periods)	u at t=0	t=6	t=12	t=24	t=36	S at t=36
0.03	22.76	0.6000	0.4998	0.4163	0.2889	0.2004	0.5887
0.05	13.51	0.6000	0.4411	0.3242	0.1752	0.0947	0.3211
0.10	6.58	0.6000	0.3189	0.1695	0.0479	0.0135	0.0527
0.20	3.11	0.6000	0.1573	0.0412	0.0028	0.0002	0.0008

Read the `lambda = 0.10` row with periods as months. A team that converges at 10% per month starts with `S = 1.2238` nats available and is down to `0.0527` nats three years later — 4.3% of what it had — **with the same people, the same brief, and, if you ask them, an improved sense of how well they work together.**

This is why u cannot be treated as a property you acquire. It is a quantity that decays and must be maintained. **`u` is a `delta` problem wearing a `u` costume**, and the maintenance regime for it — rotation, adversarial roles, deliberate outsiders, structured dissent — is the subject of Chapter 9 and is not optional infrastructure.

The correct posture toward cohesion is therefore not hostility. Cohesion does real work: it lowers `C_t` — coordination between people who trust each other is cheap — and it can raise `eta`, because people who know each other convert signals faster. **Cohesion is a genuine `C`-reducer and a genuine `u`-destroyer at the same time, and the engineering question is never "cohesion: yes or no" but "at what point does the marginal `C` saving stop covering the marginal `u` loss."** That crossover is computable in any specific system.

18.8 Groupthink, priced

The groupthink literature (Janis, 1972) named the syndrome well — illusion of unanimity, self-appointed mindguards, suppression of personal doubts, pressure on dissenters. What it did not do is price it. The arithmetic here does.

Take a group with N = 5 and u_pot = 0.60. Apply the syndrome as its two measurable effects: dissenters lose airtime (J falls from 1.00 to about 0.46, the severe-dominance profile), and private doubts are progressively abandoned rather than merely unspoken (u_pot decays at lambda = 0.10).

At onset: S = 1.2238. After twelve periods of both effects: u_pot = 0.1695, realised U = 0.46 x 0.1695 = 0.0780, N_eff = 1.31, S = ln(1.31) = 0.2716.

The group has fallen from 76% of its ceiling to 17% of it. **The illusion of unanimity is not an illusion. By the twelfth period the group genuinely is close to one instrument. Its members are reporting a true fact about the state of the group and drawing a false conclusion about what that fact means.** They read N_eff -> 1 as convergence on truth. It is convergence on each other.

This is the same error Chapter 11 calls the Parallel Wall, and it is worth stating in its most compact form, because it is the single most common misreading in institutional life: **agreement between instruments that were aligned by the process of agreeing is not evidence.**

The two-track correction: a converged group is not an emptier one There is a natural but wrong inference from the paragraph above — that the aligned group is therefore storing less. Chapters 34 and 40 compute the opposite, and this chapter must carry the correction. With a group's framing D and the world's push-back P, the bankable moment is M = min(r E, ln N) for r = D/P, while the *veridical* throughput is V = E min(r, 1). **An over-framed group banks MORE stored order than a calibrated one — it is simply not veridical stored order.** Dogma is easier to file than knowledge, and it files faster.

What the group has actually accumulated is stored error, Z_t = Q_t − Q^V_t, which obeys Z_{t+1} = (1−delta) Z_t + eta (M_t − V_t) — the cost term cancels exactly — and the bill arrives later as correction load C_t = c_0 + kappa Z_t. Take this chapter's own worked group, eta = 0.35, delta = 0.06, N = 6, c_0 = 0.42, and add *illustrative, badged* over-framing at r = 1.6 on exposure E = 0.90 with kappa = 0.03: M = 1.44, V = 0.90, so Z_inf = 0.35(0.54)/0.06 = 3.15 — **larger than the group's entire measured baseline Q_0 of 3.10** — and C_inf = 0.42 + 0.03(3.15) = 0.5145, a 22.5% rise in cost against an unchanged ceiling of eta ln N = 0.6271.

The lag is tau = 1/delta, which is why a converged committee looks strong for years and then becomes expensive without an event to blame. Cross-reference Chapter 34 §34.21, Chapter 40 §40.2, and Chapter 48, where this is the mechanism behind the doctrine-to-data pathology.

18.9 The room is part of *c*

Turn-taking equality is not purely a social variable. It is substantially a **physical** one, and this is where the sociology chapters and the space chapters of this manual meet.

From Chapter 14: in a reverberant room, reflected energy raises the noise floor against which every speaker competes, and signal-to-noise at the listener's ear determines intelligibility. A quiet, higher-pitched, or non-native speaker is disadvantaged by physics, and their contribution is therefore discounted, interrupted, or missed. The Lombard effect closes the loop: speakers raise their

voices in noise, which raises the noise. The room selects for volume and the group experiences the selection as a difference in confidence.

In a hard-surfaced room, the loudest speaker wins by acoustics, and the group will attribute the outcome to personality. Sabine's `RT60 = 0.161 V / A` is, on this reading, a `J`-equation. From Chapter 17: a table with a head assigns sightlines and speaking order before anyone opens their mouth. From Chapter 12 and Hall's proxemics (1966): conversational distance determines who is in a face-to-face relation and who is an audience member.

The cheapest available *c*-factor interventions are acoustic and geometric. They cost building work rather than culture work, they apply continuously without an administrator, and they require nobody to change. Absorption on the ceiling of a meeting room is a collective-intelligence intervention. It is odd to say and it follows directly.

Correction, and it is a real one: the last sentence follows from Postulate 14.2, and Postulate 14.2's causal claim has been tested and NOT CONFIRMED. Chapter 14 §14.13 sets out the falsification test the chain requires — acoustics move turn-taking equality, turn-taking equality is assumed to move `c` — and the worked case in Chapter 14 returned outcome (b) or a near miss of it: `N_eff_turn/n` rose from 0.604 to 0.812 with a matched untreated control room flat at 0.638 -> 0.650, while the preregistered performance measure improved only slightly and did not reach threshold. **The mediator moved decisively and the outcome did not follow.** Woolley et al. report a correlation; if turn-taking equality is a symptom of `c` rather than a driver, acoustic treatment buys the symptom.

So the claim this chapter is entitled to make is narrower than the one above, and it is the narrower claim that governs CI-9 and the Worked Example: **acoustic and geometric work is the cheapest available way to raise `J`, and `J` is a term in this chapter's model, not a proxy for `c` that has been shown to carry `c` with it.** Report such work as what is established — reduced listening effort, recovered access to the `ln N` ceiling for members the room was excluding, a measured rise in `J` — and not as a demonstrated collective-intelligence gain. That is exactly the posture §18.11 requires, and it costs nothing here, because every intervention in the Specification is defensible on its own terms whether or not the causal link survives testing.

18.10 A measurement protocol for *c*-proxies in a standing team

You cannot run Woolley's battery on your department, and if you could, the result would not transfer — their tasks were short, scored, and artificial by design. What you can do is measure the **correlates**, which is what a designer needs anyway, since the correlates are the handles.

Run this over four to six real working sessions. Nothing here requires special equipment.

Instrument 1 — participation shares (`J`). Record who speaks and for how long. A manual tick-per-turn tally is cheap but counts a short interjection equally with a long exposition; a stopwatch per speaker, or a recording segmented afterward, is better. Compute `p_i` as each member's share of total speaking time, then `H`, `E`, and `J = H / ln n`. Report `J` per session and its variance across sessions. **Report `E` alongside `J` — "we convened nine and ran four" lands where "J = 0.62" does not.**

Instrument 2 — the interruption ledger. Count interruptions by initiator and by target. A group can have acceptable `J` and still be running an interruption regime in which certain members are

systematically cut off and compensate by speaking longer when they get the floor. Report an interruption matrix, not a total.

Instrument 3 — disagreement audit (`u_pot`). Before a decision meeting, elicit each member's independent position in writing, privately, before any discussion. Measure the dispersion of positions. This gives you `u_pot` directly and it is the only instrument in this list that does. It also raises `u_pot` as a side effect by breaking the cascade before it forms — which is a nuisance for measurement and a benefit for everything else. State this contamination in your write-up; do not hide it.

Instrument 4 — error decorrelation (`u_pot`, harder and better). Where the group makes forecasts later scored, compute the correlation matrix of members' errors. Independent errors mean genuine u; correlated errors mean the members are one instrument, whatever their CVs say. The strongest available u instrument. It requires a history of scored predictions, which is why most organisations cannot run it and should start building the record now.

Instrument 5 — conversion audit (`eta`). Sample decisions from the log. For each, trace whether the objections raised in the meeting appear anywhere in the durable record — decision log, specification, changed default, altered procedure. Compute the fraction of raised objections that survive into structure. That fraction is a direct, crude, honest `eta` proxy for deliberative work.

Instrument 6 — cost tally (`C`). Person-hours in the room, plus preparation, plus rework attributable to coordination failure. Chapter 24 gives the procedure. Do not skip it: `C` decides whether any of the rest matters.

What to do with the numbers. Do *not* compute a single scalar and call it your team's *c*. You have not replicated Woolley et al. What you have is a parameter vector `(J, u_pot, eta, C, delta, N)` that plugs into the storage law and produces a testable prediction about `Q` — more useful than a scalar anyway, because a scalar tells you where you are and a parameter vector tells you which lever to pull.

And the standing warning, from Goodhart and Campbell: the moment `J` becomes a target it ceases to be a measure. A team told to equalise airtime will do so by having quiet members say filler, which raises `J` and lowers `S`, because filler is parallel to everything. Measure `J` as a diagnostic. Change the *arrangement*, not the *metric*. If you must publish a number, publish the pair `(J, u_pot)` — gaming `J` alone shows up immediately as a flat disagreement audit.

18.11 The honest limit, restated at length

It would be easy to end this chapter with the impression that collective intelligence is now an engineered quantity. It is not, and the gap between what has been shown and what has been built here should be stated plainly.

Proven. The surplus law and the storage law — mathematics, and dependent on the *chosen* design axiom S3. The affine effective-count form `g(u) = 1 + (N−1)u` is a modelling decision, not a discovery; if your system violates it the conclusions do not hold.

Measured. That in laboratory groups of two to five people, working short scorable tasks, a single factor predicted performance across tasks; that it did not track average or maximum member IQ; and that it correlated — modestly — with social sensitivity, turn-taking equality, and proportion of women.

Assumed. Everything connecting the two. Postulate 18.1's mapping. Postulate 18.2's linear form `U = J u_pot`. The extrapolation from a session to an institution. The stability of `eta`, `delta`, and `C` over an observation window long enough to estimate them.

What would falsify the chapter. Chapter 25 gives the machinery. Raising `J` while holding `u_pot`, `eta`, `C`, and `N` fixed should raise the driven term's contribution to `Y_{t+1}`; a shuffled-input control should destroy the fit. **If a matched autoregressive baseline predicts `Y_{t+1}` as well as the driven model, this chapter has told you nothing about your system and you should say so in public.**

What this chapter is for. Not to prove that *c* exists in your institution. To give you cheap, reversible, non-coercive changes — acoustic, geometric, procedural — individually defensible on their own terms, which the 2010 result gives some reason to expect will help, and which have no coercive failure mode if the expectation is wrong. A go-around that does not raise performance still costs almost nothing and still lets a quiet person speak. **The Specification below is chosen so that the downside of the postulate being false is close to zero.** That is the discipline that makes it legitimate to act on a modelling assumption at all.

The Specification

Numbered, checkable, buildable. `n` is the number of people in the group; `N` is the number of genuinely distinct knowledge categories present (Chapter 7); `J` is the equality index of §18.2.

CI-1 — Size. Deliberative groups shall be sized `n <= 8` unless a specific argument is recorded for exceeding it. Surplus is concave in `N` while communication cost grows as `n(n−1)/2` (Brooks). At `n = 8` the pair count is 28; at `n = 12` it is 66 — a 2.36x cost increase for a surplus gain of `1.7834 − 1.5790 = 0.2044` nats at `u = 0.55`, under 13%.

CI-2 — Distinct-category audit. Before convening, list the `N` genuinely distinct categories represented and name the person carrying each. Two people in the same category count once toward `N`. Record `N`, not `n`, in the meeting header. If `N < 3`, the meeting is an briefing, not a deliberation; label it correctly.

CI-3 — Participation measurement. `J` shall be measured on at least one session in four, by the Instrument 1 method, and reported alongside `E = exp(H)`.

CI-4 — Participation floor. Target `J >= 0.85` for deliberative sessions. At `n = 5` this corresponds to roughly the mild-skew profile or better. `J < 0.70` triggers a protocol review, not a conversation with the dominant speaker.

CI-5 — Silent-write-first. Every decision meeting shall open with 3-5 minutes of silent, individual, written position-taking before any discussion. Rationale: this is the only intervention in this list that raises `u_pot` realisation *and* blocks the Bikhchandani cascade *and* generates the Instrument 3 measurement, in one move.

CI-6 — Go-around before convergence. Before any decision is recorded, every member states a position. No exceptions for seniority, and no passing without an explicit "no view, and here is why."

CI-7 — Rank speaks last. The highest-status participant states their view after all others. Rationale: the cascade mechanism is asymmetric in status; a senior first-mover sets a public signal

that rationally suppresses subsequent private information.

CI-8 — Recorded minority position. Any decision log entry shall carry a field for dissent, and that field shall be filled or explicitly marked "none offered." An empty field is a measurement, not a formality: a run of empty fields is a u alarm.

CI-9 — Acoustic conformance. Deliberation rooms shall meet the Chapter 14 specification for reverberation time and background noise. Rationale: J is partly a physical variable and this is the cheapest permanent intervention available. Absorption is administered by the ceiling, not by a chair.

CI-10 — Sightline reciprocity. Seating geometry shall give every participant a direct sightline to every other (Chapter 17). No head of table. Rationale: turn-taking allocation is partly a geometric variable and geometry enforces continuously without an administrator.

CI-11 — Interruption ledger. Where J is below CI-4's floor, run the interruption matrix before changing anything. The fix for an interruption regime differs from the fix for a monologue regime and the aggregate J does not distinguish them.

CI-12 — Disagreement audit cadence. u_pot shall be measured by the Instrument 3 method at least quarterly. Rationale: u decays (§18.7); an unmeasured u is assumed constant and it is not.

CI-13 — u maintenance. Every standing group shall run at least one active u-maintenance mechanism: membership rotation, a designated adversarial role, a scheduled outside reviewer, or cross-assignment (Chapters 9 and 20). Rotate the adversarial role so it does not become a person's identity — once it does, the group discounts it and u falls back.

CI-14 — Cohesion crossover review. Annually, estimate the marginal C reduction and marginal u reduction attributable to group tenure and compare them. **Where cohesion's C saving no longer covers its u cost, rotate membership.** This is a scheduled review, not a crisis response.

CI-15 — Conversion artifact. No deliberative session shall end without a durable artifact: decision, changed default, written open question, or assigned experiment. "We had a good discussion" is an eta = 0 report (Chapter 4).

CI-16 — Objection survival audit. Quarterly, sample decisions and compute the fraction of raised objections traceable into the durable record. The target is not 1.0 — most objections should lose — it is that the fraction is *nonzero and stable*, which shows the channel exists.

CI-17 — No compositional screening. Nothing here licenses selecting individuals on any Woolley correlate. The correlates are group-level, modest, and in the proportion-of-women case explicitly mediated. Selection on them is out of scope (CANON §5.4) and is also the wrong lever: protocols are cheaper and reversible.

CI-18 — Right of abstention. Every protocol here shall be declinable. A member may decline the go-around, the written position, or the recorded dissent, and the decline shall be recorded without consequence. Rationale: CANON §5.4. **A protocol that removes the ability to decline is out of scope for this manual regardless of its measured effect.**

CI-19 — Metric hygiene. J shall never be an individual performance target, an appraisal input, or a published league table. Goodhart applies with unusual force here because J is trivially gameable by producing content-free speech, which raises J and lowers S.

CI-20 — Badge discipline. Any report using this chapter's mapping shall state that the mapping from Woolley correlates to storage-law parameters is a modelling assumption. **An unbadged postulate is a defect and is catalogued as one in Chapter 27.**

The Failure Modes

F18.1 — THE PARALLEL SUMMIT. *Clinical description.* A group of individually outstanding people, selected for excellence by the same criteria, trained in the same lineage, rewarded on the same axis. Everyone is good. `u ~ 0`. *Parameter damaged.* `u`, hence `N_eff -> 1`, hence `S -> ln 1 = 0`. *Diagnostic signature.* Fast unanimous decisions; the disagreement audit shows dispersion near zero; error-correlation matrix near 1; the phrase "we're all aligned on this" used approvingly. *Measurement that detects it.* Instrument 4, error decorrelation. Instrument 3 will do at a pinch. *Fix.* Import a practitioner from an unrelated trade (Chapter 9). Do not add another excellent person from the same lineage — the marginal surplus is `ln 1 = 0` and the marginal cost is real.

F18.2 — THE MONOLOGUE WITH WITNESSES. *Clinical description.* One speaker holds 50-80% of airtime. Others contribute assent tokens. *Parameter damaged.* `J`, hence realised `U`. At the 80% profile, `E = 2.10` of `n = 5`. *Diagnostic signature.* Meeting length uncorrelated with number of participants; the same person summarises every discussion; junior members' calendars show the meeting as optional-attendance. *Measurement.* Instrument 1. *Fix.* CI-5, CI-6, CI-7 in that order. Do not begin with a private conversation with the dominant speaker; the protocol works without requiring anyone to change and therefore has no failure mode in which they refuse.

F18.3 — THE DUEL. *Clinical description.* Two members consume the airtime in visible disagreement. The group reports healthy debate. *Parameter damaged.* `J`. Worse than F18.2 despite appearances: `E = 2.88` versus `3.32` for the single-dominant profile. *Diagnostic signature.* High interruption counts concentrated in one dyad; three or more members with airtime shares under 5%; the disagreement audit shows the *silent* members hold positions outside the range the duel spans. *Measurement.* Instrument 1 plus Instrument 2, plus Instrument 3 to expose the unspoken range. *Fix.* CI-6 go-around, sequenced so the duellists speak last.

F18.4 — THE COHESION DEATH SPIRAL. *Clinical description.* A team that has succeeded together converges. Comfort rises, `u` falls, performance follows on a lag of one to three years, and the fall is attributed to external conditions. *Parameter damaged.* `u_pot` (decay `lambda`), `J` (suppression at emission), `eta` (selective non-conversion of disconfirming input). *Diagnostic signature.* Rising self-reported team satisfaction with falling disagreement dispersion; onboarding described as "cultural fit"; departures described as "not a fit"; the recorded-dissent field empty for consecutive quarters. *Measurement.* Instrument 3 tracked over time. **The signature is the trend, not the level.** *Fix.* CI-13 and CI-14: scheduled rotation. This fix is unpopular in proportion to how much it is needed, which is the diagnostic in another form.

F18.5 — THE REVERBERANT ROOM. *Clinical description.* `J` is low and no protocol fixes it, because the cause is in the building. *Parameter damaged.* `J`, via signal-to-noise at the listener. *Diagnostic signature.* `J` varies systematically by *room* rather than by *group*; participants lean forward; the same team's `J` is materially higher in a smaller or softer room; the meeting is audibly louder at the end than the start (Lombard). *Measurement.* Measure `J` for the same group in two rooms. Then measure `RT60` (Chapter 14). *Fix.* Absorption — the cheapest permanent `c`

intervention in the manual, and a building trade rather than a management activity.

F18.6 — THE CASCADE. *Clinical description.* An early public statement, usually by the highest-status member, causes subsequent members to discount their private information and copy. Aggregation stops. *Parameter damaged.* u realised, and u_pot over repeated exposure. *Diagnostic signature.* Positions stated later in a meeting cluster more tightly than positions stated earlier; written pre-meeting positions are more dispersed than spoken ones; reversals are rare and, when they happen, total. *Measurement.* Compare Instrument 3 dispersion (written, pre-discussion) against dispersion in the spoken record. **The gap between them is the cascade, measured.** *Fix.* CI-5 and CI-7. One of the few failure modes with a genuinely complete structural fix: if positions are committed in writing before any public signal exists, the mechanism has nothing to work with.

F18.7 — GOODHART CAPTURE OF J. *Clinical description.* J becomes a target. Airtime equalises. Content does not. *Parameter damaged.* S, while the metric improves — the worst possible combination because it is invisible to the instrument. *Diagnostic signature.* J rises while the disagreement audit's dispersion is flat or falling; meeting duration rises; contributions from previously quiet members are procedural rather than substantive. *Measurement.* Always report (J, u_pot) as a pair. Divergence between them is the signature. *Fix.* Remove J from any target, appraisal, or published comparison (CI-19). Retain it as a private diagnostic.

F18.8 — THE UNBADGED *c*. *Clinical description.* A report claims the organisation has measured its collective intelligence, citing Woolley for a number computed by a method Woolley did not use, on a task type Woolley did not study, over a timescale Woolley did not observe. *Parameter damaged.* The organisation's ability to know when it is wrong, which is not a parameter in the law and is worth more than any of them. *Diagnostic signature.* A single scalar labelled c in a slide deck. No stated postulate. No falsification criterion. No control condition. *Measurement.* Read the methods section. If there isn't one, that is the finding. *Fix.* CI-20. Report the parameter vector, state the postulate, name the falsifier.

F18.9 — THE EMPTY DENOMINATOR. *Clinical description.* J is computed over people who attended, ignoring those who were relevant and not invited. The meeting scores well and the decision is still uninformed. *Parameter damaged.* N, silently. *Diagnostic signature.* J high, N low, and the same objection arriving from outside the room after every decision. *Measurement.* CI-2. Count categories, not attendees. *Fix.* Fix the invitation list before optimising the conversation. **A perfectly equal conversation among the wrong three categories is ln(1 + 2u) and no protocol raises the 2.**

Worked Example

The system

A hospital's **discharge safety review group**. Nine people meet fortnightly for ninety minutes to review discharge criteria, incidents, and proposed protocol changes. Four years old, two membership changes.

Survey (methods per Chapter 24). Distinct categories on the roster, audited per CI-2 — two consultants in the same specialty count once:

```
N = 6    acute medicine | nursing | pharmacy | physiotherapy/rehab
```

```
         data & audit    | patient-experience representative
```

ln N = ln 6 = 1.7918 nats. That is the hard per-interaction ceiling (Paper D, Cor 2.2).

Parameters estimated over eight sessions:

```
n        = 9 people
N        = 6 distinct categories
u_pot   = 0.55    Instrument 3, dispersion of written pre-discussion positions
eta      = 0.35      Instrument 5, objection survival into the durable record
delta    = 0.06     per fortnight; estimated from decay of protocol adherence
              during a 10-week gap when the group did not meet (Chapter 3 method)
C_t      = 0.42      person-hours in room + preparation + rework attributable
                   to coordination failure, in the same normalised units as Q
Q_0      = 3.10      order proxy Y: composite of documented current criteria,
                staff adherence audit score, and incident-review closure rate,
                     normalised (Chapter 24 §on proxy justification)
```

Proxy justification (Paper E §6.3, the most common measurement error). Y must not be a repackaging of U. Here Y is built from *adherence and closure* — downstream operational facts recorded by people outside the group who do not know its U — while U is estimated from written position dispersion inside the group. The two share no input. **A Y that is secretly U will fit beautifully and mean nothing.**

Turn-taking, Instrument 1, eight sessions. The two acute-medicine consultants together hold 61% of speaking time; the patient-experience representative holds 2%; physiotherapy holds 4%. Computed over the nine-way distribution, the equality index is

```
J = 0.52
```

Diagnosis at baseline

```
Realised  U   = J x u_pot = 0.52 x 0.55 = 0.2860
N_eff         = 1 + (6-1)(0.2860)       = 2.4300
S             = ln(2.4300)              = 0.8879 nats
eta S         = 0.35 x 0.8879           = 0.3108
delta Q_0     = 0.06 x 3.10             = 0.1860
C             =                           0.4200

Headroom   h  = eta ln N - delta Q - C
              = 0.35(1.7918) - 0.1860 - 0.4200
              = 0.6271 - 0.6060
              = +0.0211                  h > 0

One-step   Q_1 - Q_0 = eta S - delta Q - C
                     = 0.3108 - 0.1860 - 0.4200
                     = -0.2952            DECLINING
```

This is the diagnostic pattern that makes headroom worth computing. h > 0 says the regime is *feasible*: some interaction geometry grows order. The one-step calculation says the

current geometry does not. Confusing the two is why organisations despair when they should retune, or retune when they should stop.

Nine people convene fortnightly and produce the deliberative output of **2.43** effective participants. Order is falling at 0.295 units per period. Left alone, from Theorem 3.1:

period	Q_t (no change)
0	3.1000
4	2.0213
8	1.1791
12	0.5215
16	0.0081
17	model breaks: Q_17 = −0.1015

Theorem 3.5 said this in advance: Q_inf = (0.3108 − 0.4200)/0.06 = −1.8206, a negative equilibrium, which the theorem does not hide — it means **the regime is unsustainable or the linear model must be replaced.** At period 17 the admissibility condition C_t <= (1−delta)Q_t + eta S_t fails (0.4200 > 0.3185) and the model stops applying. In the building, that is not a mathematical event. It is the group's documented criteria quietly becoming a folder nobody opens.

Eight months of fortnightly meetings and the stored order falls to a sixth of where it started. The meetings are held. The room is booked. The minutes are taken. **The institution is running a maintenance ritual on a declining asset and recording attendance as evidence of health.**

What surplus would be required

From Proposition 4.1, a non-decreasing step requires S_t >= (delta Q_t + C_t)/eta:

```
S_req = (0.1860 + 0.4200) / 0.35 = 0.6060 / 0.35 = 1.7314 nats
ln N  = 1.7918 nats
```

S_req < ln N, so the ceiling does not rule it out — the same fact as h > 0. But invert S = ln(1 + 5U) to see what U it demands:

```
1 + 5 U = e^{1.7314} = 5.6495
U       = 0.9299
```

The group would need realised U = 0.93 when its roster carries u_pot = 0.55. Even at perfect equality, J = 1 gives U = 0.55. **The requirement is infeasible on this roster under any conversational protocol whatsoever.** The headroom gauge said feasible; the geometry says no. This is the discrepancy Chapter 6 warns about and that a single-number dashboard hides.

Intervention Stage 1 — turn-taking only

Apply CI-5, CI-6, CI-7, CI-9, CI-10: silent written positions first, full go-around, consultants speak last, ceiling absorption installed in the meeting room, seating changed from a boardroom table to a ring with no head. Measured over the next eight sessions, J rises to **0.86**.

```
U     = 0.86 x 0.55 = 0.4730
N_eff = 1 + 5(0.4730) = 3.3650
S     = ln(3.3650)   = 1.2134 nats      (up from 0.8879, +36.7%)
eta S = 0.35 x 1.2134 = 0.4247

Step  = 0.4247 - 0.1860 - 0.4200 = -0.1813      STILL DECLINING
```

The largest available conversational intervention improved the step by 0.1139 units per period and did not change its sign. Effective participants rose from 2.43 to 3.37 — a real gain, honestly measured — and the asset still shrinks, just more slowly:

period	Q_t (Stage 1 only)
0	3.1000
4	2.4375
8	1.9202
12	1.5164
16	1.2011
20	0.9549
24	0.7627

Q_inf = (0.4247 - 0.4200)/0.06 = 0.0783. The equilibrium is barely above zero. The intervention bought the group survival at near-nothing rather than collapse.

This is the result the chapter has been building toward and it is the reason the chapter is not a list of meeting tips. **A c-factor intervention operates on S. S enters the storage law multiplied by eta and competes against delta Q + C. If the cost term is large, no achievable S wins, and running a better meeting is a real improvement inside a losing budget.**

Anyone who stopped at Stage 1 would hold a genuine, measurable, publishable improvement in group process — and a department whose discharge criteria continued to decay. They would also, quite reasonably, conclude that the collective-intelligence literature does not survive contact with practice.

Intervention Stage 2 — the cost term

Audit C = 0.42. Its composition:

```
0.13   the ninety-minute session itself, nine people
0.06   preparation
0.15   a standing weekly sub-meeting duplicating the fortnightly agenda
0.08  rework: criteria re-litigated because prior decisions were not findable
```

Two moves, neither about conversation:

- **Retire the weekly sub-meeting.** It existed because the fortnightly meeting produced no findable record, so people met to re-establish what had been decided — low eta presenting as a cost. Saves 0.15.
- **Install the decision log** with an open-questions column and a required dissent field (CI-8, CI-15). Saves 0.07 of the 0.08 re-litigation cost.

```
C: 0.4200 -> 0.2000
```

Now:

```
eta S     = 0.4247       (unchanged; Stage 1 protocols retained)
delta Q_0 = 0.1860
C         = 0.2000

Headroom h = 0.6271 - 0.1860 - 0.2000 = +0.2411
Step       = 0.4247 - 0.1860 - 0.2000 = +0.0387      GROWING

S_req = (0.1860 + 0.2000)/0.35 = 1.1029  ->  U_req = (e^{1.1029} - 1)/5 = 0.4026
Realised U = 0.4730 >= 0.4026                          FEASIBLE, with margin
```

Steady state from Theorem 3.5:

```
Q_inf = (eta F(u*) - c*) / delta = (0.4247 - 0.2000)/0.06 = 3.7450
```

period	Q_t (Stage 1 + Stage 2)	h_t
0	3.1000	+0.2411
4	3.2414	+0.2326
8	3.3518	+0.2260
12	3.4380	+0.2208
16	3.5053	+0.2168
20	3.5579	+0.2136
24	3.5989	+0.2112

Contrast the three trajectories: doing nothing, the model breaks at period 17 with Q through zero; with the best available meeting protocol alone, `Q_24 = 0.7627` and falling toward 0.0783; with the protocol *and* the cost cut, `Q_24 = 3.5989` and rising toward 3.7450. And observe the convergence: Q climbs toward 3.7450 and slows as it goes, because h_t shrinks as Q_t grows. **Success consumes its own headroom** (Chapter 6). This is Banach's fixed point theorem doing its work: `delta = 0.06 > 0` makes the map a contraction, so the unique fixed point exists and the iteration arrives at it. Book One's Node 10 FIXEDPOINT is Paper E Theorem 3.5, and here it is with numbers on it.

What the group should watch next

Three things, in order of when they will bite.

The cohesion clock. `u_pot = 0.55` was measured, not assumed, and it will decay (§18.7). At `lambda = 0.05` per fortnight it halves in 13.5 periods — about a year. If u_pot falls to 0.40, then `U = 0.86 x 0.40 = 0.344`, `S = ln(2.72) = 1.0006`, `eta S = 0.3502`, and the step at `Q = 3.60` becomes `0.3502 - 0.2160 - 0.2000 = −0.0658` — negative again, with nothing visibly changed and everyone getting on better than ever. Re-run Instrument 3 quarterly.

The maintenance trap. As Q rises toward 3.745, `delta Q` rises from 0.1860 to 0.2247. The budget tightens by 0.0387 purely from success. This group is small enough that it does not bite; Chapter 23 is about the case where it does.

The proxy. The moment any of Y's three components becomes a target, Campbell's Law applies and the proxy detaches from the thing. Watch the adherence audit — it is the most gameable.

The falsification statement

Preregistered per Chapter 25, before Stage 2 was implemented:

```
MODEL      Y_{t+1} = a Y_t + b F(U_t) − Chat_t + e_t ,  0 <= a <= 1, b >= 0
FIT        on periods −16..0 (pre-intervention), then FROZEN
PREDICT    periods 1..24 held out
BASELINES  (i) persistence: Y_{t+1} = Y_t
           (ii) direct-input: Y_{t+1} = a Y_t + b U_t   (raw u, no log)
           (iii) matched autoregressive with no interaction term
       (iv) SHUFFLED-INPUT CONTROL: the same F(U_t) series randomly permuted across period

FALSIFIED IF
  1. the shuffled-input control predicts held-out Y as well as the ordered series;
  2. the matched autoregressive baseline is not beaten;
  3. the fitted a lies outside [0,1] or the fitted b is negative.
```

The shuffled-input control is the one most often omitted and the one that decides this case. A department that improved its meetings, cut two commitments, installed a decision log, and got better may simply be a department that received attention. The shuffle asks whether the *timing* of surplus predicts the *timing* of order. If it does not, the arithmetic in this chapter described a coincidence in the vocabulary of a theorem, and the honest report says so.

Chapter 19 . Contagion, Thresholds, and Critical Mass

The Law

PROVEN (Papers A and E). The storage law and the surplus law, unchanged:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t ,     S_t = ln(1 + (N−
1)U_t) ,  0 <= S_t <= ln N
h_t     = eta ln N − delta Q_t − C_t
```

This chapter adds no mathematics to that. What it adds is a theory of **how a practice arrives in a system at all** — because every intervention in this manual is a practice that somebody has to adopt, and an intervention nobody adopts changes no parameter.

MEASURED — the tipping point. Centola, Becker, Brackbill & Baronchelli, *Science* 360 (2018). In controlled online experiments, groups established an arbitrary naming convention through repeated pairwise interaction. The experimenters then introduced a **committed minority** — confederates who consistently advanced an alternative convention and never deviated. Below a critical size the minority was absorbed and the incumbent convention held. Above it, **the alternative convention displaced the established one across the whole group.** The measured tipping point clustered around **25% of the population.**

That is the finding, stated exactly, and the precision matters because it is routinely repeated in looser and false forms. It is not "25% of people change the world." It is: *in an experimentally constructed system, with a coordination convention, with unanimously committed confederates, the proportion at which the incumbent convention gave way was near one quarter.*

MEASURED — complex contagion. Centola & Macy (*AJS*, 2007) drew the distinction that makes the rest of this chapter work. A **simple contagion** transmits on a single contact: a virus, a rumour, a piece of trivia. A **complex contagion** requires **multiple, independent, reinforcing contacts** before an individual adopts: costly behaviours, risky behaviours, behaviours whose value depends on others adopting, behaviours that carry social risk if you are the only one doing them.

The consequence is structural and it inverts the standard intuition. For simple contagion, random long-range ties are enormously valuable — they compress path length and reach distant regions fast (Watts & Strogatz, *Nature*, 1998). For complex contagion, a long tie delivers **one** contact, and one contact is below threshold. **Clustered topology beats random long ties**, because clustering is precisely the property that delivers a second and third contact from within a person's own neighbourhood. Centola's experimental study (*Science*, 2010) tested this directly, comparing behaviour spread on clustered-lattice against random network structures.

MEASURED — thresholds. Granovetter (*AJS*, 1978) modelled collective behaviour with individual thresholds: person *i* acts when the number or fraction of others already acting reaches `T_i`. The aggregate outcome is determined by the *distribution* of thresholds, not by average disposition, and the mapping from distribution to outcome is **discontinuous** — arbitrarily small changes in one person's threshold can change the outcome from nothing to everything.

POSTULATED, BADGED.

```
POSTULATE 19.1 (adoption as a parameter driver). BADGED: modelling assumption.
  A practice adopted by fraction a of a population changes the storage-law
  parameters roughly in proportion to a:
```

```
    eta(a) = eta_0 + (eta_1 - eta_0) a
    C(a)   = C_0   + (C_1   - C_0)   a
  Linearity in a is chosen for tractability, not derived, and should be replaced
   by a measured curve wherever one exists.

POSTULATE 19.2 (field transfer). BADGED, and this is the big one.
   The ~25% figure and the clustering advantage are treated in this manual as
   ORDER-OF-MAGNITUDE DESIGN GUIDANCE for real organisations. This is an
   extrapolation from controlled online experiments to institutional settings
   with hierarchy, memory, exit options, material stakes, and pre-existing
   relationships. NONE of those were present in the experiments.
```

HONEST LIMIT. The 2018 result comes from **controlled online experiments** with anonymous participants, an artificial coordination task, no organisational hierarchy, no history, no material consequence, and a minority that was **perfectly committed by construction** — confederates do not waver, get tired, get promoted, or leave. Real committed minorities do all four. **Field transfer of the 25% number is a modelling assumption, not a demonstrated result.** Treat 25% as a sanity-check magnitude that tells you whether your seeding plan is off by a factor of ten. Do not treat it as a threshold you can compute to two decimal places and then declare a mandate on. That misuse is common, it is not supported by the paper, and it is catalogued as a failure mode below.

CANON §5.4 applies with unusual force in this chapter. A method that describes how practices spread is also a method that could be used to make practices spread against people's judgement. Everything specified here is restricted to arrangements people can decline. §19.11 is not an appendix to that principle; it is a load-bearing section.

The Engineering

19.1 Why the intuitive seeding plan is wrong

Ask any competent manager to spread a new practice across a 220-person organisation. The plan is almost always the same: identify a champion in each of the twelve teams, brief them, give them material, and let it spread outward.

The plan maximises **reach**. It puts a seed within one hop of the largest possible number of people for the smallest possible number of seeds. It is exactly the right plan for a simple contagion, and it is the wrong plan for almost every practice an organisation actually wants to spread.

Here is the mechanism of the failure, and it is worth stating before any arithmetic. A person encountering a new practice through exactly one colleague faces a specific and rational inference problem. **One adopter is indistinguishable from an enthusiast.** The observer cannot tell whether they are seeing a signal about the practice or a signal about the person. A second, independent adopter — someone whose opinion is not derivative of the first — resolves the ambiguity. That is what "reinforcing contact" means, and it is not irrationality. It is correct Bayesian caution about a correlated source.

Reach delivers first contacts. Clustering delivers second contacts. Almost everything worth spreading needs the second.

19.2 The reinforcement requirement, made numerical

Model the requirement crudely and honestly: a person adopts when at least a of their k neighbours have already adopted. Take k = 8 and a = 2 — the mildest possible complex contagion, requiring just one more contact than a virus. Under a *uniformly random* placement of seeds at population fraction s, the probability that a given non-seed has at least one and at least two seeded neighbours is:

seed fraction s	P(>= 1 seeded nbr)	P(>= 2 seeded nbrs)	ratio
0.05	0.3366	0.0572	0.170
0.10	0.5695	0.1869	0.328
0.114	0.6203	0.2294	0.370
0.15	0.7275	0.3428	0.471
0.20	0.8322	0.4967	0.597
0.25	0.8999	0.6329	0.703
0.30	0.9424	0.7447	0.790

Read the s = 0.10 row. Sprinkling seeds across 10% of the population puts **57%** of everyone within one hop of the new practice, which looks like magnificent coverage and will be reported as such. It puts only **19%** within reinforcing range. **The gap between 57% and 19% is the entire distance between a launch that looks successful in week one and a practice that has vanished by month three.**

And note the shape: raising a from 1 to 2 — the smallest possible step into complex contagion — cuts effective coverage by a factor of three at low seed density. The penalty is worst exactly where budgets force you to operate.

19.3 Clustering: what it is and what it buys

The random-placement arithmetic above assumes the *worst* case: seeds scattered so that no two land adjacent. Real topology is not random, and the property that matters is the **clustering coefficient** — the probability that two of your neighbours are also neighbours of each other.

Two structures with identical mean degree k = 8 on n = 220 people:

```
RING LATTICE (each node joined to its 8 nearest neighbours)
  clustering  C = 3(k−2) / (4(k−1)) = 3(6)/(4(7)) = 18/28 = 0.6429
  mean path   ~ n / (2k) = 220/16 = 13.75 hops

RANDOM GRAPH (same mean degree)
  clustering  C ~ z/n = 8/220 = 0.0364
  mean path   ~ ln n / ln z = 5.394/2.079 = 2.594 hops
```

The lattice is **17.7x** more clustered. The random graph has paths **5.3x** shorter.

For a simple contagion, the random graph wins overwhelmingly: reach everyone in under three hops instead of fourteen. For a complex contagion, the lattice wins, and the mechanism is visible in the numbers: in the lattice, an adopter's neighbour has on average 8 x 0.6429 = 5.14

neighbours who are also neighbours of that adopter. In the random graph the same figure is 8 x 0.0364 = 0.29. **The lattice hands each newly-exposed person five potential reinforcing contacts. The random graph hands them one third of one.**

Watts & Strogatz's rewiring parameter p interpolates between them, and clustering falls approximately as C(p) = C(0)(1−p)^3:

rewiring p	clustering C	expected reinforcing contacts among 8 neighbours
0.00	0.6429	5.14
0.05	0.5512	4.41
0.10	0.4686	3.75
0.20	0.3291	2.63
0.50	0.0804	0.64
1.00	0.0000	0.00

The small-world regime — the one that makes information travel beautifully — is the regime in which behaviour stops travelling. At p = 0.10 path length has already collapsed toward the random-graph value while clustering retains most of its lattice value; that is the celebrated small-world sweet spot for *information*. Push to p = 0.50 and you have a superb rumour network and a system in which no costly practice can establish, because a person's neighbours no longer know each other and therefore cannot corroborate.

This has a direct and unwelcome implication for the modern organisation. Every tool that flattens communication — the all-hands channel, the company-wide forum, the cross-functional matrix, the open-plan floor that puts everyone within earshot of everyone — **raises p**. Each is defended, correctly, on information grounds. Each degrades the substrate that costly practices need. **An organisation optimised end to end for information flow has, without noticing, optimised itself against behaviour change.**

19.4 Granovetter thresholds and the discontinuous aggregate

Granovetter's model is the simplest interesting model in sociology and it has the most counter-intuitive output.

Give each of 100 people a threshold T_i: the number of others who must already be acting before person *i* acts. Start with nobody acting and iterate r_{t+1} = |{ i : T_i <= r_t }| to a fixed point.

Case A — thresholds 0, 1, 2, …, 99. One person with threshold 0 acts. That makes one actor, so the person with threshold 1 acts. Now two, so threshold 2 acts. Simulated: the cascade runs 100 steps and terminates with **all 100 acting.**

Case B — identical to A, except that the single person with threshold 1 instead has threshold 2. Person 0 acts. Nobody has threshold 1. Nobody else acts. Simulated: the cascade terminates immediately with **1 acting.**

Case C — identical to A, except nobody has threshold 0. Nothing happens at all. **0 acting.**

One person's threshold moved by one unit is the difference between a movement of one hundred and a movement of one. The average threshold is essentially unchanged. The distribution of dispositions is essentially unchanged. Every survey instrument you could run on this population before the event would return the same answer in all three cases.

Three engineering consequences follow.

First: aggregate sentiment is not a predictor. "Most people are broadly in favour" tells you nothing about whether anything will happen, because the outcome is set by the *shape of the low tail* of the threshold distribution, not by its mean. A population with a favourable average and a gap at $T = 1$ is inert.

Second: the binding constraint is usually a gap, not a shortage. In Case B the population is overwhelmingly willing; the cascade fails because one particular rung of the ladder is missing. The correct intervention is to find and fill the gap, which is cheap, not to persuade the population, which is expensive and in Case B unnecessary.

Third: the zero-threshold seed is structurally required. Case C shows a population where every single person would act given a single predecessor, and where nothing happens forever, because nobody will go first. **This is why "someone has to be the first to try it" is a technical statement about a fixed point and not an inspirational one.** The role of a designed pilot is to manufacture a $T = 0$ actor whose going-first is institutionally sanctioned rather than personally risky.

19.5 Threshold heterogeneity is the property you must measure

Granovetter's real lesson is that the *variance* of the threshold distribution matters as much as its location, and it is the quantity nobody collects.

In an organisation, thresholds are not random. They are structured by position:

population	typical threshold driver	typical T
people who proposed the practice	none; already committed	0
people with slack and low personal risk	curiosity	low
people whose work is measured on old terms	measurable cost of switching	high
people whose expertise is embedded in the old practice	status loss	very high
people with one foot out of the organisation	no return on adoption	effectively infinite

The fourth row is the important one and it is routinely misread as resistance-to-change in the psychological sense. **A person whose accumulated skill is a stock of Q denominated in the old practice has a rationally high threshold, because adoption writes down their asset.** They are not being difficult. They are doing correct arithmetic on a different balance sheet. Chapter 20 is entirely about this case, and its answer — build the lower valley rather than push — is the threshold-lowering move.

Practical instrument. Before seeding, ask a sample: *"How many colleagues in your immediate team would need to be using this before you would try it?"* The answers are self-reported and noisy and they are still worth more than a satisfaction survey, because they estimate the shape of the low tail. What you are looking for is not the mean. **You are looking for whether there is anybody at `T <= 1`, and whether the counts at `T = 2` and `T = 3` are large enough to keep the ladder unbroken.**

19.6 Percolation: why change looks like nothing and then everything

The threshold model explains *whether* a cascade runs. Percolation explains *what it looks like from inside while it is running.*

Take a random graph on `n` nodes with mean degree `z`. The fraction `S` of nodes in the giant connected component satisfies `S = 1 - e^{-zS}`, solved by iteration:

mean degree z	giant component S	nodes of 220
0.5	0.0000	0.0
0.8	0.0000	0.0
0.9	0.0000	0.0
1.0	0.0010	0.2
1.1	0.1761	38.7
1.2	0.3137	69.0
1.5	0.5828	128.2
2.0	0.7968	175.3
3.0	0.9405	206.9
4.0	0.9802	215.6

Between `z = 1.0` and `z = 1.1` — a 10% change in connectivity — the largest connected group goes from essentially zero to thirty-nine people. Below the threshold there is no giant component at any size; above it, one appears and grows fast. This is a phase transition, and the population on either side of it experiences it very differently from how it appears in the table.

The lived consequence, and it is the reason so many good practices are killed:

Below threshold, every observation says failure, and the observation is correct. Adoption clusters stay small, isolated, and stagnant. They do not merge. Reports say "some enthusiasm in two teams, no broader traction." A reasonable executive cancels the programme.

Just above threshold, the observations say the same thing for a while. The giant component exists but is still small in absolute terms; the qualitative reports are indistinguishable from the below-threshold case. Then it grows, and the growth is fast, and everyone involved retrospectively constructs a narrative about the moment when people finally came around.

There was no such moment. **The parameter crossed a threshold and the visible behaviour lagged it.**

This produces the single most expensive error in organisational change: **cancelling at the point of maximum apparent failure, which is frequently the point just before the transition.** The defence is not faith. It is instrumentation:

```
DO NOT MEASURE:  total number of adopters
DO MEASURE:      size of the LARGEST CONNECTED CLUSTER of adopters
                 number of distinct clusters
                 number of edges BETWEEN clusters
```

Total adopters rises smoothly and tells you nothing about which side of the transition you are on. **Largest-cluster size is the order parameter, and it is the only number that distinguishes "slow" from "stuck."** Two hundred isolated adopters is a failed programme. Forty adopters in one connected component is a programme about to work.

19.7 The 25% result, and what it is for

Now the tipping-point number, and the discipline for using it.

Centola et al. (2018) is valuable to an engineer for a reason slightly different from the one it is usually cited for. Its value is not the specific quarter. Its value is that it **demonstrated experimentally that a tipping point exists at all** — that the relationship between committed minority size and convention change is not gradual, that a minority below the critical size is absorbed with no residue, and that a minority above it flips the whole population. Before that result, tipping points in social conventions were a theoretical prediction. After it, they are a measured phenomenon in at least one carefully constructed setting.

The practical value of the number is as an **order-of-magnitude check**:

organisation scope	population	25%	committed people needed
whole firm	220	55.0	55
one division	60	15.0	15
one unit	30	7.5	8
one team	12	3.0	3
one squad	6	1.5	2

Stare at that table, because it contains the most useful design move in the chapter.

SHRINK THE DENOMINATOR. A committed minority of 55 people is a political campaign requiring budget, sponsorship, and eighteen months. A committed minority of 3 is a team lunch. The threshold is a *proportion*, so **the cheapest way to get above it is not to recruit more committed people but to narrow the population within which the convention must hold.**

This is why the correct unit of change is almost never the organisation. Establish the convention in one team of 12 with 3 committed members. That team is now, internally, past its tipping point, and the convention there is self-sustaining rather than sponsored. It is also now a **cluster** in the sense of §19.3 — a dense region where anyone adjacent receives multiple reinforcing contacts. The second team requires 3 more committed people, and some of them are now available, because they are the first team's members who work partly with the second.

Sequential local tipping is cheaper than simultaneous global tipping by a wide margin, and it is the only version that produces clusters as a side effect rather than requiring them as a precondition.

19.8 Cluster, do not spray: the calculation

Concretely, a 220-person firm with a budget for 25 people's time to establish a new practice. Assume mean degree k = 8 and a 2-of-8 reinforcement rule.

PLAN A — SPRAY. Distribute the 25 across the organisation, roughly two per team, maximising reach. Seed fraction s = 25/220 = 0.1136.

```
P(a given non-seed has exactly 0 seeded neighbours) = 0.3810
P(exactly 1)                                        = 0.3907
P(exactly 2)                                        = 0.1753
P(exactly 3)                                        = 0.0450

P(>= 1 seeded neighbour)  = 0.6203     <- the number that gets reported
P(>= 2 seeded neighbours) = 0.2294     <- the number that matters

Expected non-seeds within reinforcing range = 195 x 0.2294 = 44.5 people
```

Note the second row: **39% of the population has exactly one seeded neighbour.** That is the worst-populated cell in the table. Those 76 people have encountered the practice, formed an impression, and their impression is *"one person on my team is doing an odd new thing."* They are not neutral afterwards. **A below-threshold exposure is not a free trial; it spends the practice's novelty and returns a null result, and the second attempt to introduce the same practice to the same person starts from a worse position than the first.**

PLAN B — CLUSTER. Put all 25 into a single 30-person unit. Local seed fraction inside that unit is 25/30 = 0.8333.

```
P(>= 2 seeded neighbours) for the 5 non-seeds in the unit = 1.0000
Expected activatable inside the unit = 5.00 of 5
Expected activatable outside the unit ~ 0 in the first period
```

Plan A's headline number — 44.5 people in reinforcing range — is nearly nine times Plan B's five. Any reasonable reading of a one-period comparison prefers Plan A.

And Plan A is the wrong choice, for three reasons that a one-period comparison cannot see.

First, the 44.5 are at the exact margin. They have precisely two contacts and the rule requires two. Any attrition among the 25 seeds — one person on holiday, one reassigned, one quietly stopping because their two teammates did not join — drops many of them below threshold. Plan A's coverage is brittle in the exact way that a two-of-eight rule is brittle.

Second, Plan A's adopters are not connected to each other. Per §19.6, the order parameter is largest-cluster size, and Plan A maximises the number of clusters while minimising the size of each. It is engineered, precisely, to sit below the percolation threshold.

Third, Plan B's unit is past its own tipping point. 25 of 30 is 83%, far above 25%. The convention inside that unit is not a programme; it is how things are done. It survives the sponsor losing interest, which Plan A does not. And it is now a genuine cluster: anyone in an adjacent unit who works with two or three people from the seeded unit receives multiple independent reinforcing contacts, which is exactly the condition Plan A could not manufacture anywhere.

Plan A optimises the first-period metric. Plan B optimises the order parameter. The first-

period metric is what gets reported and the order parameter is what determines the outcome.

The correct seeding procedure, then:

1. Pick the smallest population within which the practice makes sense standing alone. It must be able to work if NOBODY outside adopts. If it cannot, pick a different practice.
2. Seed that population past its own tipping point. Use 25% as the floor, not the target; the experimental figure came from perfectly committed confederates and yours are people.
3. Let it stabilise until the practice survives the removal of its sponsor. That is the test. Remove yourself for a month and see.
4. Seed the adjacent population using members of the first as the reinforcing contacts. Do not send an evangelist. Send two people who work there anyway.
5. Repeat outward along real working ties, not along the org chart.

Step 5 deserves its own emphasis. **Seed along the graph of who actually works with whom, which is not the graph of who reports to whom.** Chapter 20 gives the structural reason: Conway's Law says the system mirrors the communication structure, and the communication structure is the real graph, not the published one. Seeding along the org chart is seeding along a diagram.

19.9 What "committed" must mean, and what it must not

The 2018 experiment's minority was committed by construction — confederates instructed never to deviate. Reproducing that property in an organisation is where this chapter could become something this manual refuses to be.

Two mechanisms produce apparent commitment, and they are not equivalent.

Mechanism 1: people who have used the practice, find it better, and continue using it because it is better. Their commitment is a report on the practice. Their persistence under social pressure is evidence. This is the mechanism the chapter is written for.

Mechanism 2: people who are required to display the practice, or whose standing depends on displaying it. Their commitment is a report on the incentive. Their persistence is evidence about the incentive and nothing else.

The two are behaviourally identical from outside and completely different in what they mean. Mechanism 2 produces a minority that looks committed, hits 25%, and flips the convention — **and the convention it establishes carries no information, because the adopters' agreement was manufactured rather than aggregated.** This is Chapter 11's Parallel Wall arriving through a new door: u collapses to zero across the adopter population because every adopter's position derives from the same source.

The arithmetic is unforgiving here. A practice adopted by mechanism 2 across a population produces `u -> 0` among adopters, `N_eff -> 1`, and `S -> ln 1 = 0`. **You have successfully spread a practice and set the surplus of the population that adopted it to zero.** The spread was the objective and the objective was wrong.

Hence:

Mandating adoption defeats the purpose of adoption. Not for ethical reasons alone, though those are sufficient under CANON §5.4, but because the mechanism that makes a spread practice

valuable — independent people independently finding it better — is exactly the mechanism a mandate destroys. The mandate delivers the observable and destroys the referent. It is Goodhart's Law operating on adoption itself.

The design test, applied to every seeding plan:

```
THE DECLINE TEST
  Can a person in the seeded population decline to adopt, indefinitely,
  with no consequence to their standing, assignments, or evaluation?

  If NO -> this is not seeding. It is a mandate wearing seeding's vocabulary,
           it is out of scope for this manual (CANON §5.4), and the resulting
           adoption statistics do not mean what they appear to mean.

  If YES -> the adoption rate is a measurement. Proceed, and record the
            declines, because the decline rate is your best available signal
            about the practice.
```

Recorded declines are the most valuable data a seeding programme generates and the data most often discarded. A practice with 60% adoption and 40% articulate declines is better understood than one with 100% compliance, and it is far more likely to be worth keeping.

19.10 Coupling contagion to the storage law

Contagion determines a, the adopted fraction. Postulate 19.1 turns a into parameters. The storage law turns parameters into Q. The chain runs:

```
topology + seeding + thresholds  ->  a_t  ->  (eta_t, C_t)  ->  Q_t
```

Two consequences worth stating before the Worked Example demonstrates them.

First, the adoption dip. Most practices worth adopting cost something during transition: the people using the new way and the people using the old way must interoperate, and interoperation across two conventions is more expensive than either convention alone. In the parameter language, C rises before eta does. **A programme that is working will show Q falling for several periods after launch.** This is not a warning sign. It is the expected signature. It is also, reliably, when the programme is cancelled.

Second, partial adoption can be strictly worse than zero adoption. If a stalls at a level where the dual-convention cost is fully incurred and the efficiency gain is only partly realised, the system sits at a step worse than it started. Two conventions running in parallel is the worst of the three available states. **This is the strongest available argument for the cluster-first method: a saturated small unit reaches full local adoption and skips the dual-convention regime entirely, whereas a sprayed population is engineered to live in it.**

19.11 The honest limits, stated in full

On the 25%. It is one experiment, in one paradigm, online, with anonymous participants, on an arbitrary naming convention with no stakes, with perfectly committed confederates. Real conventions have material stakes, real minorities waver, real populations have hierarchy and memory and exit options. The experimental design's strength — clean control — is exactly what limits its

transfer. **Use 25% to check whether your plan is off by a factor of ten. Do not use it as a target and never use it as a justification for a mandate.**

On complex contagion. The distinction is well-supported and the clustering advantage was demonstrated experimentally by Centola (2010). But whether a *particular* practice in *your* organisation is complex, and what its a-of-k rule is, is unmeasured in your case. The `2-of-8` rule used throughout this chapter is an illustration, not a finding.

On the network model. Ring lattices and Erdos-Renyi random graphs are analytical conveniences. Real organisational networks are neither: they are clustered, degree-heterogeneous, and partly hierarchical. Barabási & Albert (*Science*, 1999) showed preferential attachment produces scale-free degree distributions with hubs; Albert, Jeong & Barabási (*Nature*, 2000) showed such networks are robust to random failure and fragile to targeted hub attack. **The corresponding seeding hazard: a practice that spread through a hub is a practice with a single point of failure, and hubs move jobs.** Seeding through clusters produces redundancy where seeding through hubs produces speed.

On the postulates. 19.1's linearity in a is a convenience. 19.2's field transfer is the whole bridge and it is not proved. If your measured `Q` trajectory does not track `a`, the postulate is wrong for your system and the honest response is to say so and report the parameter estimates anyway, because a failed prediction with published parameters is worth more than a successful narrative.

The Specification

CN-1 — Classify the contagion. Before seeding, record whether the practice is simple (adoptable on one exposure, no social risk, no dependence on others adopting) or complex. If complex, record the estimated a-of-k rule and the evidence for it. **Every practice in this manual — decision logs, go-arounds, rotation, shared instruments — is complex. Assume complex unless argued otherwise.**

CN-2 — Map the real graph. Seed along the graph of who works with whom, evidenced by shared artifacts, shared calendars, or observed collaboration. **The org chart is not the graph.** Record the method used to build the map.

CN-3 — Measure clustering. Estimate the clustering coefficient of the seeding region. If it is below roughly 0.2, seeding will require a materially larger committed fraction, because reinforcing contacts must be manufactured rather than harvested.

CN-4 — Pick the smallest viable population. The seeded population must be one within which the practice delivers value even if nobody outside adopts. If it cannot stand alone, the practice is mis-scoped, not under-resourced.

CN-5 — Seed past the local tipping point. Target a committed fraction of at least 25% **of the chosen small population**, treated as a floor. Justification: Centola et al. (2018), with Postulate 19.2's field-transfer caveat recorded alongside the number.

CN-6 — Never spray. Distributing a fixed seeding budget uniformly across a large population is prohibited without a recorded argument that the contagion is simple. It maximises the first-period reported metric and minimises the order parameter.

CN-7 — Threshold survey. Before seeding, sample the target population with: *"How many colleagues in your immediate team would need to be doing this before you tried it?"* Report the **distribution**, especially the counts at `T = 0, 1, 2, 3`. A gap in the low tail is a stop condition, not a reason to persuade harder.

CN-8 — Manufacture the zero-threshold actor. Provide at least one person whose going-first is institutionally sanctioned and personally costless. **A population in which everyone would follow and nobody will lead is stable at zero forever** (§19.4, Case C).

CN-9 — Order parameter, not headcount. Programme reporting shall track largest connected adopter cluster, number of clusters, and inter-cluster edges. Total adopter count may be reported but shall not be the primary metric.

CN-10 — Declare the dip. Before launch, state in writing that `Q` (or its proxy) is expected to fall for a stated number of periods, with the predicted trough depth and recovery period. **A dip predicted in advance is evidence; a dip explained afterward is an excuse.**

CN-11 — Cancellation criteria, pre-registered. State before launch what observation would justify stopping, in terms of the order parameter and the predicted trough. Absent this, programmes are cancelled at the point of maximum apparent failure, which is often immediately before the transition (§19.6).

CN-12 — The decline test. No seeding programme shall proceed unless a member of the seeded population can decline indefinitely with no consequence to standing, assignment, or evaluation (CANON §5.4).

CN-13 — Record declines. Declines shall be recorded with reasons and reported alongside adoption. A programme reporting adoption without declines is reporting half the measurement.

CN-14 — No mandate. Adoption shall not be required, incentivised by evaluation, or used as a performance signal. Rationale is not only ethical: mandated adoption drives `u –> 0` across the adopter population and sets the surplus of the spread practice to `ln 1 = 0` (§19.9).

CN-15 — Sponsor-removal test. Before expanding beyond the first seeded population, the sponsor shall withdraw active support for one full period. If the practice decays, it has not established and expansion is premature.

CN-16 — Expand along working ties. The second population shall be seeded by people who already work in it, not by visiting advocates. Two embedded practitioners beat one evangelist, by exactly the arithmetic of §19.2.

CN-17 — Hub caution. Where seeding through a high-degree individual, record the single-point-of- failure risk explicitly and seed at least one cluster that does not depend on that person (Albert, Jeong & Barabási, 2000).

CN-18 — Avoid the dual-convention plateau. Do not plan for sustained partial adoption across a population. Either saturate a small population or do not start; the parallel-convention regime can be worse than no adoption at all (§19.10).

CN-19 — Watch p. Where an organisation is simultaneously introducing communication tools that flatten structure — global channels, matrix reporting, open floors — record that these raise the effective rewiring probability and degrade the substrate for behaviour change, even as they improve information flow (§19.3).

CN-20 — Badge discipline. Any use of the 25% figure shall carry the statement that it derives from controlled online experiments and that field transfer is a modelling assumption. **An unbadged 25% is a defect** (Chapter 27).

The Failure Modes

F19.1 — THE SPRAY. *Clinical description.* A fixed seeding budget distributed uniformly to maximise reach. High first-contact coverage, near-zero reinforcement coverage. *Parameter damaged.* `eta`, via a that never rises; and the practice's future viability, because below-threshold exposure spends novelty. *Diagnostic signature.* Many adopter clusters of size 1-2; largest-cluster size flat; adoption plateaus at roughly the seeded fraction and does not grow. *Measurement.* Largest connected adopter cluster over time. If it equals the largest seed group after three periods, nothing spread. *Fix.* Consolidate. Withdraw seeds from most locations and concentrate them until one population is past its local tipping point. This looks like retreat and is the only move that works.

F19.2 — THE HUB BET. *Clinical description.* Seeding through a small number of well-connected individuals. Fast initial spread, no redundancy. *Parameter damaged.* Robustness of `a`. Also `u`, because everything downstream of a single hub is correlated with that hub's judgement. *Diagnostic signature.* Adoption traceable to one or two people; the practice's justification, when sampled across adopters, is reported in near-identical language. *Measurement.* Trace each adopter's first two exposures. If they converge on one person, that is the finding. *Fix.* CN-17. Establish one independent cluster before the hub leaves. **Test by asking what happens if that person is promoted next quarter, and then act as though they will be.**

F19.3 — CANCELLATION AT THE TROUGH. *Clinical description.* A programme is stopped during the expected adoption dip, at the point of maximum apparent failure, which is frequently just before the transition. *Parameter damaged.* All of them, plus the organisation's future capacity to seed anything, since each cancelled programme raises thresholds for the next. *Diagnostic signature.* Cancellation decision cites falling output; no pre-registered trough prediction exists; largest-cluster size was rising at the time of cancellation. *Measurement.* CN-10 and CN-11 in advance. Retrospectively, plot largest-cluster size against the cancellation date. *Fix.* Pre-register the dip and the cancellation criteria. **The purpose of pre-registration is to make the decision when you are calm and to bind yourself when you are not.**

F19.4 — THE MANDATE. *Clinical description.* Adoption required or tied to evaluation. Adoption statistics reach 100%. *Parameter damaged.* `u -> 0` across adopters, so `N_eff -> 1` and the spread practice carries `S = 0`. Also the organisation's ability to learn that the practice is bad. *Diagnostic signature.* Adoption near 100%; declines near zero; usage-depth metrics far below usage-breadth metrics; the practice performed in the presence of observers. *Measurement.* Sample depth of use rather than presence of use. Sample declines — if there are none, that is the signature, not a success. *Fix.* Out of scope for repair by this manual; the correct action is to remove the mandate and re-measure. CN-12 and CN-14 exist to prevent it.

F19.5 — THE DUAL-CONVENTION PLATEAU. *Clinical description.* Adoption stalls partway. Both conventions run. Interoperation cost is fully incurred; efficiency gain is partly realised. The system is worse than before it started. *Parameter damaged.* `C`, badly. `eta`, partly. *Diagnostic signature.* Adoption flat for three or more periods at 20-60%; rising complaints about translation, duplication, and "which system do I use"; rework rising. *Measurement.* Track `C` directly, not just

a. The dual-convention cost is visible as rework. *Fix.* Choose. Either saturate the population or withdraw the practice cleanly. **The plateau is the one state that is worse than either endpoint, and staying in it out of unwillingness to reverse is the most common way this failure becomes permanent.**

F19.6 — THE MISSING RUNG. *Clinical description.* A population overwhelmingly favourable in survey, in which nothing spreads, because the threshold distribution has a gap in its low tail. *Parameter damaged.* a, absolutely. *Diagnostic signature.* High stated support, zero adoption; the threshold survey shows nobody at `T <= 1`, or a gap at `T = 2`. *Measurement.* CN-7 threshold survey. **This failure is invisible to every sentiment instrument and obvious to the threshold instrument.** *Fix.* Fill the gap. Recruit or sanction one or two specific people at the missing rung. Cost: trivial. Persuading the population instead: expensive and unnecessary.

F19.7 — THE REWIRED ORGANISATION. *Clinical description.* An organisation that has systematically flattened its communication structure finds that information travels instantly and no practice ever establishes. *Parameter damaged.* Effective clustering, hence reinforcement probability, hence a. *Diagnostic signature.* Excellent information reach; every initiative known to everyone within a day; nothing sticks; a history of practices that "everyone knows about" and nobody uses. *Measurement.* Estimate clustering of the real working graph (CN-3). Compare with the ratio in §19.3. *Fix.* Deliberately re-cluster for adoption purposes: stable small teams, co-assignment, durable pairings. **This is a real trade against information flow and must be recorded as one, not sold as free.**

F19.8 — THE UNBADGED QUARTER. *Clinical description.* "Research shows you only need 25% to change a culture," cited in support of a programme, with no mention of the paradigm, the confederates, or the artificiality of the convention. *Parameter damaged.* The organisation's ability to know when it is wrong. *Diagnostic signature.* The number appears without a method. It is used to justify a target rather than to sanity-check a magnitude. *Measurement.* Ask what the 25% was 25% *of*, in the study. *Fix.* CN-20.

F19.9 — REACH REPORTED AS TRACTION. *Clinical description.* A launch reports the fraction of the organisation "reached" or "aware." The figure is high, real, and irrelevant. *Parameter damaged.* None directly. It damages the decision process, which then damages everything. *Diagnostic signature.* Reporting is in units of awareness, attendance, or page views; the word "engagement" appears; no cluster structure is reported. *Measurement.* Convert the reported reach into `P(>= 2 reinforcing contacts)` using the §19.2 table. The gap between the two numbers is the finding. *Fix.* CN-9. **Report the order parameter or report nothing.**

Worked Example

The system

A **220-person professional services firm**, twelve teams, chronically losing decisions. The same questions are re-litigated across engagements; nobody can find why a prior call was made; senior staff spend a measurable fraction of their week reconstructing context.

The proposed practice is a **decision log**: for every decision above a stated threshold, a short structured record of the decision, its alternatives, the dissent, and the open questions. It is a

classic eta device (Chapter 4) and a classic complex contagion — it is only worth writing if others read it, and it is socially costly to be the only person writing them.

Survey (Chapter 24 method).

```
n        = 220 people, 12 teams
N        = 8 distinct practice categories (audit, tax, advisory, technology,
           legal, data, sector specialists, client-relationship)
ln N     = 2.0794 nats
u        = 0.50    from disagreement audit on sampled engagement decisions
S        = ln(1 + 7 x 0.50) = ln(4.50) = 1.5041 nats
delta    = 0.05    per month; estimated from decay of decision-recall accuracy
                    over a 6-month window with no logging
C_0      = 0.55    coordination overhead, normalised
eta_0    = 0.30    objection-and-rationale survival into any findable record
Q_0      = 4.00    order proxy Y: composite of (i) fraction of sampled decisions
                    whose rationale is reconstructible from records, (ii) re-
litigation
                   rate, (iii) median time to answer "why did we do it this way"
```

Proxy justification (Paper E §6.3). Y is built from reconstructibility, re-litigation rate, and lookup latency — all recorded by sampling engagement files, none of them derived from the disagreement audit that produces U. **A Y that is secretly U fits beautifully and means nothing;** these share no input.

Estimated post-adoption parameters, from a two-team pre-pilot measured over four months:

```
eta_1 = 0.52    rationale survival rises sharply once a structured record exists
C_1   = 0.46     falls: less reconstruction work, fewer re-litigations
              (net of the cost of writing the logs, which is real and included)
```

Baseline diagnosis

```
eta_0 S     = 0.30 x 1.5041 = 0.4512
delta Q_0   = 0.05 x 4.00   = 0.2000
C_0         =                 0.5500

h_0  = eta_0 ln N - delta Q_0 - C_0 = 0.30(2.0794) - 0.2000 - 0.5500
     = 0.6238 - 0.7500 = -0.1262            h < 0

step = 0.4512 - 0.2000 - 0.5500 = -0.2988   DECLINING
```

h_0 < 0, and Corollary 3.4 is categorical about what that means: no choice of interaction geometry can prevent order from decreasing. No amount of better meetings, better recruitment, or higher u fixes this firm, because even at the ceiling S = ln N = 2.0794 the input eta_0 ln N = 0.6238 is less than delta Q + C = 0.7500. **The binding constraint is eta and C, and the diagnosis is available before anyone is consulted about anything.**

Post-adoption, if the practice fully lands:

```
eta_1 S = 0.52 x 1.5041 = 0.7821
h_1     = 0.52(2.0794) - 0.2000 - 0.4600 = 1.0813 - 0.6600 = +0.4213
```

```
step_1  = 0.7821 − 0.2000 − 0.4600 = +0.1221                GROWING
Q_inf   = (0.7821 − 0.4600)/0.05 = 6.4424
```

So the practice is worth having. The remaining question is entirely a contagion question: does it arrive?

The seeding decision

Budget: 25 people's time. Two plans, priced in §19.8's arithmetic.

Plan A, spray: two seeds per team. `P(>= 2 reinforcing contacts) = 0.2294`; 44.5 people in range; zero teams past their local tipping point; largest adopter cluster = 2.

Plan B, cluster: saturate one 30-person team. Local committed fraction 83%, far above the 25% floor; largest adopter cluster = 25; the practice's value is realised locally within that team, because the team's own decisions become reconstructible regardless of the other eleven teams (CN-4).

Plan B is chosen. Note explicitly what is given up: for roughly two months, 195 people see no change at all, and the programme's reportable reach is 14%. **CN-10 requires stating this in advance, because a reach of 14% at month two will otherwise be read as failure.**

The adoption trajectory

Adoption then proceeds outward from the seeded team along real working ties (CN-16). Modelled logistically with `r = 0.55` per month, `K = 220`, `a_0 = 25/220`, and coupled to the storage law by Postulate 19.1:

```
eta(a) = 0.30 + 0.22 a          C(a) = 0.55 − 0.09 a
Q_{t+1} = Q_t + [ eta(a_t) S − delta Q_t − C(a_t) ]
```

month	adopters	share a	eta	C	Q_t	h_t	step
0	25.0	0.1136	0.3250	0.5398	4.0000	-0.0640	-0.2509
1	37.2	0.1690	0.3372	0.5348	3.7491	-0.0211	-0.2151
2	54.2	0.2463	0.3542	0.5278	3.5340	+0.0320	-0.1718
3	76.6	0.3484	0.3766	0.5186	3.3622	+0.0965	-0.1203
4	104.1	0.4732	0.4041	0.5074	3.2419	+0.1708	-0.0617
5	134.3	0.6103	0.4343	0.4951	3.1802	+0.2490	-0.0009
6	163.1	0.7412	0.4631	0.4833	**3.1793**	+0.3206	+0.0542
7	186.3	0.8467	0.4863	0.4738	3.2335	+0.3757	+0.0959
8	202.0	0.9181	0.5020	0.4674	3.3294	+0.4100	+0.1212
9	211.1	0.9594	0.5111	0.4637	3.4506	+0.4266	+0.1325
10	215.8	0.9808	0.5158	0.4617	3.5831	+0.4317	+0.1349
12	219.1	0.9960	0.5191	0.4604	3.8505	+0.4266	+0.1279
14	219.8	0.9992	0.5198	0.4601	4.1009	+0.4158	+0.1167

Read the `Q_t` column. **It falls for six months, bottoms at 3.1793 — 20.5% below where it started — and does not recover to its starting value until month 14.**

This is the adoption dip (§19.10), computed rather than asserted, and it is the reason CN-10 exists. Month five is the maximum danger point: the step is `−0.0009`, essentially flat, `Q` is at its worst,

61% of the firm has adopted, and every visible operational indicator is at its lowest reading since before launch. A cancellation at month five would be entirely defensible on the evidence available at month five and would destroy a programme that was one month from turning.

Note also the h_t column, which tells a different and earlier story. **Headroom crosses zero at month two, four months before Q turns.** The headroom gauge is the leading indicator; the order proxy is the lagging one. A firm watching h_t sees the regime become sustainable in month two. A firm watching Q_t sees nothing good until month six. **This is precisely why Chapter 6 insists on computing h, and it is the strongest practical argument in the manual for doing so.**

Sizing check against the tipping point

At month 2, a = 0.2463, essentially the Centola figure. And indeed month 2 is where h turns positive. **This coincidence is a coincidence** and must be labelled as such: the logistic curve's parameters were chosen from the pre-pilot, and h crossing zero at a = 0.246 reflects the eta(a) and C(a) slopes, not any tipping-point mechanism. Two arithmetic facts landing on the same month is not evidence that one caused the other. **Reporting it as confirmation of the 25% threshold would be exactly the misuse this chapter's honesty clause exists to prevent.**

The genuine tipping-point calculation is the one in the seeding decision: 25 committed people inside a 30-person team is 83%, far above the floor, which is why that team's practice held after the sponsor withdrew (CN-15, tested at month 3 with a one-month withdrawal — logging continued at 92% of its prior rate).

Falsification statement, pre-registered

```
MODEL      Y_{t+1} = a Y_t + b F(U_t) − Chat_t + e_t ,  0 <= a <= 1,  b >= 0
FIT        on months −12..0, then FROZEN
PREDICT    months 1..18 held out

PRE-REGISTERED PREDICTIONS
 P1  Y falls for 5–7 months, trough between 3.05 and 3.30, recovery to Y_0 by month 12–16.
  P2  h crosses zero between months 1 and 3.
  P3  largest connected adopter cluster exceeds 60 by month 4.
   P4   the seeded team sustains >=85% logging rate through a one-month sponsor withdrawal.

BASELINES  (i) persistence  (ii) direct-input, raw U with no log
           (iii) matched autoregressive, no interaction term
           (iv) SHUFFLED-INPUT CONTROL

FALSIFIED IF
 1. the shuffled-input control predicts held-out Y as well as the ordered series;
  2. the matched autoregressive baseline is not beaten;
  3. fitted a outside [0,1] or fitted b negative;
  4. the trough falls outside [2.8, 3.5] or occurs outside months 4–9.
```

Prediction P1 is the one that makes this preregistration worth the paper. Predicting that a

programme will *work* is cheap — every programme predicts that. Predicting the depth and timing of the hole it digs first, before digging it, is a claim that can embarrass you in a specific month. That is what makes it evidence rather than advocacy, and it is why the honest version of this chapter's method is not "seed in clusters and it will spread" but **"seed in clusters, state in advance what the failure will look like on its way to working, and let the calendar decide."**

Chapter 20 . The Silo as Local Minimum

The Law

PROVEN (Paper A, Theorem 4.1 — the decomposition that defines this chapter). For a pure block vector, the non-parallelism splits into two independent sources:

```
u = u_cross + u_div        u_cross = beta^2         (BETWEEN categories)
                         u_div   = alpha^2 sin^2(gamma)   (WITHIN a category)

max( f(u_cross), f(u_div) )  <=  f(u)  <=  f(u_cross) + f(u_div)
```

The upper bound is subadditivity on the parameter u (Theorem 3.4). **No synergy beyond additivity is promised and honest design does not assume one.**

PROVEN (Paper E, §3 and Corollary 3.4).

`Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t ,   h_t = eta ln N − delta Q_t − C_t`

and if `C_t > eta ln N − delta Q_t`, **no choice of interaction geometry prevents order from decreasing.**

This chapter's central structural claim, stated in the proven vocabulary:

> **A silo is an arrangement in which `u_cross` exists on the roster and is realised at approximately zero, while `C` for any interaction that would realise it is multiplied by the cost of crossing the boundary.**

That is the whole diagnosis, and it names why silos are unusually damaging: they attack the surplus term and the cost term at once. Most failure modes damage one parameter. **The silo is a u-destroyer and a C-multiplier simultaneously, which is why interventions that address only one of the two reliably make things worse.**

CITED — the structural cause. Conway's Law (1967): organisations which design systems produce designs which are copies of the communication structures of those organisations. **BADGE: Conway's Law is an observation and a widely-supported design heuristic, not a theorem.** It is used here in its contrapositive, which is the operationally useful direction: *if you want the system to have a seam somewhere, put a communication boundary there; if you do not want a seam, do not put one there.* Silos in the output are the shadow of boundaries in the org.

CITED — the contact literature. Allport (1954) proposed that contact between groups reduces prejudice **under conditions** — equal status within the situation, common goals, cooperation rather than competition, and institutional support. Pettigrew & Tropp's meta-analysis (2006), covering more than 500 studies, found contact typically associated with reduced prejudice, with the effect generally stronger where Allport's conditions held. The engineering lesson transfers directly and is usually ignored: **contact is not the intervention. Contact under specified conditions is the intervention, and contact without those conditions is a cost with no benefit.** §20.9 prices exactly that case.

POSTULATED, BADGED — the comfort landscape.

```
POSTULATE 20.1 (the landscape). BADGED: an ANALOGY, not a theorem and not a measurement.
  A practitioner's position is modelled as a point on a scalar potential V(x) over a space
  of practices, where lower V = more comfortable. Behaviour is modelled as descent with
```

```
  agitation: people move downhill readily, uphill rarely and briefly.

POSTULATE 20.2 (the escape-rate analogy). BADGED: an ANALOGY BORROWED FROM PHYSICS.
  The rate of moving from one valley to another is modelled as
      rate ~ exp( - barrier / T )
  where T indexes willingness to tolerate discomfort. This is an Arrhenius/Kramers form
  used here for its SHAPE, not because any thermodynamic law has been established.
  CANON §5.1 forbids the claim that the surplus law proves anything about physics, and
  nothing here claims the converse either. It is a diagram with numbers on it.
```

THE SAFETY PROPERTY — the reason this chapter's method is the one this manual recommends.

The method specified below works by **removing resistance** rather than by applying force. It lowers the ground somewhere else and lets people walk there if they wish.

A method that works by removing resistance has no coercive failure mode. Its worst outcome is that someone declines to move.

That is not a soft statement. It is a hard design property, and it is checkable: for every instrument in the Specification, ask what happens to a person who does not want it. For a rotation programme that is offered, the answer is *they do not rotate.* For a shared instrument that is built, the answer is *they do not open it.* For a mandate, a reorganisation, or a performance metric tied to cross-functional behaviour, the answer is *they comply or lose standing*, which is CANON §5.4 territory and out of scope for this manual. **The landscape method is preferred here not only because the arithmetic below favours it but because its failure mode is a shrug.**

The Engineering

20.1 What a silo is, precisely

A silo is not a group of people who dislike other groups. That version is a morality tale and it generates useless interventions — mixers, offsites, values statements — that fail for reasons the tale cannot explain.

A silo is a **local minimum on a comfort landscape**, with two defining geometric properties:

DEEP. Every direction out is uphill. Inside the silo, a practitioner is fluent, competent, correct, and quick. Their tools work, their vocabulary is understood on first use, their judgements are recognised as expert, their stored `Q` — accumulated skill, reputation, working relationships, tacit knowledge — is fully denominated in the local currency. **Training and reward carved the valley, and they carved it deliberately and well.** A twelve-year specialist is deep in a valley that twelve years of correct incentives dug for them. The depth is not a pathology. **It is the accumulated `Q` of a person who did exactly what the institution asked.**

NARROW. It admits one vocabulary. A question posed in a neighbouring dialect does not merely take longer to answer; it frequently cannot be recognised as a question. Where a term is shared across boundaries but carries different referents — *risk*, *model*, *validation*, *requirement*, *done* — the narrowness is worse than an unknown word, because both sides believe they have understood. **The silo's boundary is not a wall. It is a bad translation that neither side knows is happening.**

Those two properties distinguish an actual silo from mere specialisation:

```
SPECIALISATION      people are expert in different things and can still transact
SILO              people are expert in different things and transaction across the
                boundary costs so much that it stops happening, after which the
                divergence in vocabulary accelerates because nothing corrects it
```

The second clause is the important one. **A silo is self-deepening.** Reduced contact means reduced correction of divergent vocabulary, which raises translation cost, which further reduces contact. A positive feedback loop that runs whether or not anyone intends it.

20.2 The landscape, made numerical

Take the model landscape (Postulate 20.1, badged as analogy) as a double well:

```
V(x) = 16 B (x^2)(1−x)^2  −  D x

  x = 0   the silo
  x = 1   the neighbouring practice
  B = 1.0 comfort units, the barrier height with no tilt
  D       the DEPTH ADVANTAGE given to the neighbouring valley
```

`D = 0` is the untouched situation: two valleys of equal depth with a barrier of 1.0 between them. `D > 0` is what happens when you make the neighbouring practice genuinely more comfortable — better tools, cheaper translation, a shared instrument that answers questions from both sides.

Computed by scanning the potential:

D	silo min x_A	ridge x_T	ridge height	new min x_B	new min depth	**barrier out of the silo**	depth advantage
0.0	0.000	0.500	0.0000	1.000	0.0000	**1.0000**	0.0000
0.2	0.006	0.487	0.9013	1.006	-0.2006	**0.9019**	0.2000
0.4	0.013	0.475	0.8050	1.012	-0.4024	**0.8076**	0.3999
0.6	0.020	0.463	0.7113	1.018	-0.6054	**0.7171**	0.5996
0.8	0.027	0.450	0.6201	1.024	-0.8095	**0.6307**	0.7990
1.0	0.035	0.436	0.5315	1.028	-1.0148	**0.5483**	0.9980
1.5	0.056	0.403	0.3216	1.042	-1.5324	**0.3609**	1.4931
2.0	0.081	0.365	0.1295	1.054	-2.0562	**0.2029**	1.9828
2.5	0.115	0.321	-0.0424	1.065	-2.5858	**0.0794**	2.4641
3.0	0.174	0.250	-0.1875	1.075	-3.1210	**0.0040**	2.9295

Two structural facts in this table matter more than the individual numbers.

First: deepening the neighbouring valley lowers the barrier out of the silo. The ridge does not merely move; its height above the silo floor falls from 1.0000 to 0.0040 as `D` goes from 0 to 3. You did nothing to the silo. You did not persuade anyone. **The uphill got shorter because the downhill on the other side got longer.**

Second: the ridge migrates toward the silo. x_T moves from 0.500 to 0.250. The point at which a person crosses over into "this is now easier" arrives earlier in their journey. In practice: with a good shared instrument, a specialist does not have to become half-fluent in the other domain before the crossing pays. They have to become a quarter-fluent.

20.3 Why exhortation fails, in two independent ways

Exhortation is the standard intervention: leadership states that silos are bad, that collaboration is a value, that everyone must think cross-functionally. It fails reliably. The landscape model explains why in two distinct ways, and both are worth having because they suggest different corrections.

FAILURE MODE 1 — exhortation is a tilt that is removed when the speech ends.

A sustained push of magnitude P toward x = 1 is mathematically identical to a tilt D = P. So a sufficiently forceful exhortation *does* lower the barrier while it is being applied:

```
P = 0.2  ->  barrier while the push is on = 0.9019
P = 0.5  ->  barrier while the push is on = 0.7619
P = 1.0  ->  barrier while the push is on = 0.5483
when the push ends -> barrier returns to 1.0000
```

The barrier returns. And crucially, so does the relative depth of the valleys: the neighbouring valley is no lower than it was, so anyone who did cross during the push window finds themselves in a place that is not more comfortable than where they came from. **They walk back, and they walk back carrying evidence that the initiative did not work, which raises their threshold against the next one** (Chapter 19, F19.3).

FAILURE MODE 2 — exhortation raises agitation, and agitation is undirected.

The other reading of exhortation is that it does not change the landscape at all; it raises T, the willingness to tolerate discomfort. Under Postulate 20.2 that does raise the crossing rate. It also raises the crossing rate *back*, by exactly the same factor, and that symmetry is the whole problem.

Compare the two routes. **Route A: lower the far valley (raise D). Route B: raise agitation (raise T), landscape untouched.** With T = 0.35 and barrier B = 1.0 at baseline:

route	forward barrier	reverse barrier	forward rate	reverse rate	**forward / reverse**
A: D = 0.0	1.0000	1.0000	0.057433	0.057433	**1.000**
A: D = 0.5	0.7619	1.2616	0.113409	0.027198	**4.170**
A: D = 1.0	0.5483	1.5463	0.208789	0.012060	**17.313**
A: D = 1.5	0.3609	1.8540	0.356566	0.005006	**71.231**
A: D = 2.0	0.2029	2.1857	0.560129	0.001940	**288.655**
B: T = 0.35	1.0000	1.0000	0.057433	0.057433	**1.000**
B: T = 0.50	1.0000	1.0000	0.135335	0.135335	**1.000**
B: T = 0.70	1.0000	1.0000	0.239651	0.239651	**1.000**
B: T = 1.00	1.0000	1.0000	0.367879	0.367879	**1.000**

Route B's forward-to-reverse ratio is exactly 1.000 at every level of agitation, forever. No amount of exhortation produces net directed movement, because it has not made anywhere better

than anywhere else. What it produces is *more movement in both directions* — which in an organisation looks like churn, restructuring, initiative fatigue, and people trying things and going back. It is not that exhortation does nothing. **It is that exhortation does the wrong thing energetically: it agitates a system whose landscape it has not changed.**

Route A, meanwhile, is asymmetric by construction. At `D = 1.0`, the forward barrier has fallen from 1.0000 to 0.5483 **and the reverse barrier has risen from 1.0000 to 1.5463.** The crossing is seventeen times more likely forward than back. That asymmetry is what "the change stuck" means, mechanically.

There is one more comparison worth computing, because it is the trap:

```
halving the barrier at T = 0.35    ->  exp(−0.5/0.35) = 0.239651
doubling T with barrier at 1.0     ->  exp(−1.0/0.70) = 0.239651
```

The two routes produce identical forward rates. A manager watching only the forward rate — how many people are trying the new thing — cannot distinguish them. The difference appears only in the reverse rate, which is to say **only after the initiative's attention moves on.** This is why exhortation looks like it is working for exactly as long as anyone is measuring.

20.4 The method: build a lower valley next door

The method is not to push. It is to build a lower valley next door — to make the cross-domain view the more comfortable place to stand.

This is the operational content of the chapter, and each word of it is doing work.

Build. Something is constructed. There is a deliverable, it has a cost, and someone maintains it. This is not a communication exercise.

A lower valley. Genuinely lower. If the cross-domain position is not actually more comfortable — faster to work in, cheaper to be understood in, better instrumented — then `D = 0`, the barrier stays at 1.0000, and the method has not been applied. **The most common failure is to declare a lower valley without building one.**

Next door. Adjacent to the existing practice, reachable by an incremental step. A valley on the far side of two ridges is not an option, however deep. The x_T migration in §20.2 is the useful version of this: a good instrument moves the crossing point from 0.500 to 0.250, so the required step shortens.

The more comfortable place to stand. Comfort here is Book One's operational sense — the reading on the gradient, not a synonym for ease. It includes being understood on first attempt, having your question answered without a translation round-trip, and not losing status for asking.

And the acid test, which is the safety property in operational form: if nobody moves, nothing bad happened. You built an instrument that is available and unused. That is a wasted cost and it is a *bounded* one, and it damaged nobody. Compare the failure of a reorganisation.

20.5 What actually lowers the valley

Five instruments, ordered by cost-effectiveness in the systems this manual is written for. Each is a `D`-raiser and most are simultaneously `C`-reducers, which is the point.

INSTRUMENT 1 — SHARED INSTRUMENTATION. One measurement apparatus, one dataset, one dashboard, one definition of the quantity in dispute, visible to both sides and trusted by both.

This is the highest-leverage item in the chapter and the most underused. Its mechanism is precise: **a shared instrument removes the translation step from the most frequent class of cross-boundary interaction — the argument about what is actually happening.** Two domains that disagree about a fact must first negotiate whose measurement counts, and that negotiation is the bulk of the boundary cost. Remove it and the cost per cross-boundary exchange falls sharply.

Requirements, drawn from the ones that fail: *readable* by both sides without specialist tooling; *trusted* by both, which usually means neither side owns it; *current*, because a stale shared instrument is worse than none since it produces confident wrong agreement; and answering questions that *both* sides actually ask, which means specified jointly rather than by whoever commissioned it.

INSTRUMENT 2 — BOUNDARY OBJECTS. Star & Griesemer (1989) named the class: artifacts that are plastic enough to be adapted to local needs on each side and robust enough to maintain a common identity across sides. A specimen, a map, a standardised form, a shared repository, a prototype.

The property that makes a boundary object work is counter-intuitive: **it does not require the two sides to agree on what it means.** A drawing can be a manufacturing instruction to one group and a regulatory artifact to another, and coordination succeeds without either side adopting the other's interpretation. **This is why boundary objects lower the valley without requiring anyone to move into it — the highest possible score on the safety property.**

Design note: a boundary object standardised into a single canonical interpretation has stopped being a boundary object and become a form. Forms are `C`-reducers and not `D`-raisers, because they require one side to translate into the other's frame before writing.

INSTRUMENT 3 — THE TRANSLATOR ROLE. A person whose explicit job is to carry meaning across the boundary, with the standing to do so and the time to do it.

This works and it has a specific structural hazard: the translator becomes a hub, and per Chapter 19's discussion of Albert, Jeong & Barabási (2000), a structure that depends on a hub is robust to random failure and fragile to that hub's departure. **A translator who is doing the job well makes themselves indispensable, which means the boundary cost has been paid rather than removed.** The correct use is transitional: the translator's success condition is that the boundary cost falls to the point where they are no longer needed, which requires them to be building Instruments 1 and 2 while they translate, not merely translating.

INSTRUMENT 4 — ROTATION. People spend structured time working inside the neighbouring practice — not visiting, not observing, working.

Rotation is a direct `u_cross` realiser: it converts roster difference into contacted difference. Its dosage is computable (§20.9). Its hazards: a rotation short enough to be comfortable is short enough to be tourism, and one long enough to build fluency is long enough that the sending team resists releasing the person. It must be **offered**, not assigned, or the safety property is lost.

INSTRUMENT 5 — CO-LOCATED CROSS-FUNCTION. Seating people from different domains within conversational range, on shared work rather than adjacent work.

The physical chapters supply the specification: Hall's proxemics (Chapter 12) gives the distances

at which spontaneous exchange occurs; Whyte's direct-observation method (Chapters 12 and 15) gives the means to check whether it is occurring rather than assuming it; Gehl's optional activity is the proxy for whether the arrangement is working. **Co-location without shared work produces proximity without contact, and proximity without contact is a `C` increase with no `u` gain** — the exact configuration §20.9 prices as strictly harmful.

20.6 Conway's Law: the structural cause

Conway (1967): a system's structure mirrors the communication structure of the organisation that built it. Badge: an observation and a heuristic, not a theorem.

For this chapter the useful reading is causal and directional. **The silo is not primarily a psychological fact about the people in it. It is the organisational shadow of a boundary that was drawn somewhere, usually for a reason that made sense at the time.** Reporting lines, budget envelopes, physical floors, tooling choices, and professional accreditation all draw boundaries, and each drawn boundary produces a corresponding seam in whatever the organisation produces.

Three practical consequences.

First, the silo is downstream of a structural decision and is not repairable upstream of it. If the budget is allocated by function, cross-function work has no budget line, and there is no amount of collaborative culture that creates one. Look for the boundary before looking at the behaviour.

Second, boundaries are legitimate and necessary. Chapter 5's arithmetic requires them: `C` grows as `n(n−1)/2`, and an organisation with no internal boundaries pays the full quadratic. **Boundaries are a `C`-control device and they work.** The question is never "remove the boundaries" but "is this boundary where the seam should be, and is its crossing cost proportionate."

Third — the deepest point, and it is Alexander's. "A City is Not a Tree" (1965) argues that living structures are **semilattices**, in which sets overlap, rather than **trees**, in which they nest without overlap. An organisation chart is a tree by construction. **A tree assigns each person exactly one category, which is a statement that `u_cross` will be realised only through the tree's root — that is, through management.** The semilattice permits overlapping membership: a person belongs to a function *and* a product *and* a guild *and* a project, and the overlaps are precisely where `u_cross` gets realised without routing through the top.

Hierarchy is the organisational form that drives `u` toward zero (Chapter 8), because it is the form that permits exactly one path between any two people and charges the full boundary cost for it. Semilattice structure — overlapping, non-nesting memberships — is the structural version of building a lower valley: it does not persuade anyone to cross a boundary; it arranges for there to be fewer boundaries between any two given people.

20.7 The double indictment, priced

Now the arithmetic that makes this chapter more than a metaphor. Consider a 60-person organisation in two silos of 30.

The `u` side. Use Theorem 4.1 with `N = 6` categories, and take measured-style values: the between-silo non-parallelism `u_cross = beta^2 = 0.55`, and the within-silo dissent `u_div =`

```
alpha^2 sin^2(gamma) = 0.0990.

u_cross = 0.5500      f(u_cross) = ln(1 + 5(0.5500)) = 1.3218
u_div   = 0.0990      f(u_div)   = ln(1 + 5(0.0990)) = 0.4021
u       = 0.6490      f(u)       = ln(1 + 5(0.6490)) = 1.4457

Theorem 4.1 bound check:
  max(1.3218, 0.4021) = 1.3218  <=  1.4457  <=  1.3218 + 0.4021 = 1.7239    HOLDS
```

Now the silo condition: `u_cross` is on the roster and is **realised at approximately zero**, because the interactions that would realise it do not occur.

```
Available  S = f(u_cross + u_div) = 1.4457 nats
Realised   S = f(u_div only)      = 0.4021 nats
Forgone                           = 1.0436 nats per interaction  =  72.2%

For reference, the hard ceiling is ln N = ln 6 = 1.7918 nats.
```

The organisation is discarding 72% of the surplus its own roster can supply, and the discarded portion is the cross-category portion — the part that is expensive to hire and impossible to substitute. Note also that this is not a diversity-of-hiring failure. The difference was purchased. It sits in the building. It is not in contact.

The `C` side. Brooks' arithmetic on the same 60 people:

```
paths(30) = 30 x 29 / 2 = 435
paths(60) = 60 x 59 / 2 = 1770

within-silo pairs = 2 x 435 = 870
cross-silo pairs  = 30 x 30 = 900
check: 870 + 900 = 1770 = paths(60)                                  CONSISTENT
```

More than half of all pairs in this organisation — 900 of 1770 — are cross-boundary pairs. Model the cost as `C = c_in x (within pairs in contact) + m x c_in x (cross pairs in contact)`, where `m` is the **boundary-cost multiplier**: how many times more expensive a cross-boundary exchange is than a within-boundary one. With `c_in = 0.00012` and all pairs in contact:

`m`	`C`	cross-boundary share of `C`
1	0.7080	50.8%
2	1.0680	67.4%
4	1.7880	80.5%
8	3.2280	89.2%
16	6.1080	94.3%

At `m = 8` — a plausible figure where two domains share terms with different referents and no shared instrument exists — **89% of the organisation's coordination cost is boundary-crossing cost.**

Now put the two sides together, because this is where the silo's peculiar viciousness shows up:

The silo suppresses the interactions that would realise `u_cross`, and the reason it suppresses them is that `m` makes them expensive. So the organisation is simultaneously forgoing 72% of its available surplus AND, to the extent that cross-boundary contact happens at all, paying a multiplied price for it. Every attempt to fix the first by mandating more contact makes the second worse. This is why reorganisations aimed at collaboration so often produce a measurably worse system: **they raise contact without lowering `m`, and the arithmetic of that combination is unambiguous.**

20.8 The critical boundary-cost multiplier

Let `rho` be the fraction of cross-boundary pairs actually in contact. Realised non-parallelism is modelled (Postulate: linear, badged) as `u(rho) = rho x u_cross + u_div`, and cost as `C(rho) = c_in x 870 + m x c_in x rho x 900`.

The step is `eta ln(1 + 5 u(rho)) − delta Q − C(rho)`. Differentiate with respect to `rho` and set to zero:

```
d(step)/d(rho) = eta (N−1) u_cross / (1 + (N−1) u(rho))  −  m c_in x 900

m* = eta (N−1) u_cross / [ (1 + (N−1) u(rho)) x c_in x 900 ]
```

`m*` is the **critical boundary-cost multiplier**: more contact improves the system only while `m < m*`.

`rho`	`u(rho)`	`m*` at `eta = 0.34`	`m*` at `eta = 0.46`
0.00	0.0990	5.79	7.84
0.10	0.1540	4.89	6.62
0.20	0.2090	4.23	5.73
0.30	0.2640	3.73	5.05
0.45	0.3465	3.17	4.29
0.60	0.4290	2.75	3.72
0.80	0.5390	2.34	3.17
1.00	0.6490	2.04	2.76

The load-bearing statement of this chapter's arithmetic:

> **Whether cross-boundary contact helps or harms is determined by `m`, the boundary cost — not by the willingness, goodwill, or collaborative disposition of the people crossing.**

If `m > m*`, then increasing contact makes the system monotonically worse, and it does so no matter how enthusiastically people participate. Enthusiasm does not appear in the expression. This is the precise sense in which **collaboration is not a virtue problem.**

Notice too that `m*` **falls as `rho` rises**: the first cross-boundary contacts are the ones that pay, because surplus is concave in `u` while cost is linear in `rho`. Full integration is rarely optimal. **The optimum is partial contact at low boundary cost, and the two policies that fail are "integrate everything" and "integrate nothing."**

And notice that m* **rises with eta**: at eta = 0.46 instead of 0.34, the tolerable boundary cost rises by about 35% at every rho. **Raising the conversion rate makes contact affordable that was not affordable before.** Since a shared instrument raises eta *and* lowers m at the same time, it moves both sides of the inequality in the favourable direction, which is why it ranks first in §20.5.

20.9 What Allport's conditions are, in this vocabulary

Pettigrew & Tropp's meta-analysis found the contact effect generally stronger under Allport's conditions. Translated into the parameters:

Allport condition	what it does in the model
equal status within the situation	prevents one side's frame becoming canonical, preserving u_cross rather than converting it to u = 0 by absorption
common goals	ensures the interaction has an output, so eta > 0
cooperation, not competition	keeps C from acquiring a conflict term
institutional support	supplies the standing and time that make m payable, i.e. lowers m

Read the fourth row against §20.8. **Institutional support is not encouragement. It is the thing that lowers m below m*.** A leadership statement supporting collaboration that comes with no budget line, no time allocation, no shared instrument, and no change to how work is assigned lowers m by nothing, leaves m > m*, and therefore makes the system worse in exact proportion to how successfully it increases contact.

Contact without institutional support is not a weak version of the intervention. It is the intervention with its sign flipped.

20.10 The honest limits

The landscape is an analogy. Postulates 20.1 and 20.2 are badged and they are not measurements. Nobody has measured a comfort potential for a professional practice, B = 1.0 is an arbitrary unit, and T = 0.35 is a chosen illustration. What the analogy supplies is a *shape* — that barriers matter exponentially and tilts matter asymmetrically — and the shape is what drives the design recommendation. **If the shape is wrong, the recommendation is wrong, and the way to find out is to measure m and rho directly, which does not require the analogy at all.**

The u_cross = 0.55 and u_div = 0.099 figures are illustrative. They are internally consistent with Theorem 4.1 and they are not measurements of any real organisation. Chapter 24 gives the procedure for obtaining real ones.

The m model is crude. Treating boundary cost as a single scalar multiplier over all cross-boundary pairs is a simplification; real boundary costs vary enormously by pair and by task. **The m* result survives the crudeness in its qualitative form — there is a critical cost above which contact harms — but the specific numbers do not transfer.**

Conway's Law is a heuristic. Widely supported, not proved, and there are organisations whose outputs do not mirror their communication structure. Use it to generate hypotheses about where a seam came from, not to explain seams you have not investigated.

Linear `u(rho)` is a modelling choice. Realised difference probably does not rise linearly in contact fraction; the first contacts likely realise more per unit than the later ones, which would make the concavity argument in §20.8 stronger, not weaker. The linear form is the conservative one.

The Specification

SL-1 — Locate the boundary before diagnosing the behaviour. Record what draws the boundary: reporting line, budget envelope, floor plan, toolchain, accreditation, or vocabulary. **A silo is downstream of a structural decision** (§20.6) and cannot be repaired upstream of it.

SL-2 — Measure `u_cross` and `u_div` separately (Chapter 10). Report both. An organisation rich in one and starved of the other requires a different intervention from one starved of both.

SL-3 — Measure `rho`, the fraction of cross-boundary pairs actually in contact. Evidence: shared artifacts, shared meetings, co-authored outputs, message graphs. Not stated intentions.

SL-4 — Estimate `m`, the boundary-cost multiplier. Method: time a representative sample of within-boundary and cross-boundary exchanges of comparable substance, from question asked to answer usable. `m` is the ratio. Crude and sufficient.

SL-5 — Compute `m*` before increasing contact. Using the §20.8 expression with your measured `eta`, `u_cross`, and current `rho`. **If `m > m*`, increasing contact is prohibited until `m` is reduced.** This is the single most important rule in the chapter and it is routinely violated by reorganisations.

SL-6 — Lower `m` before raising `rho`. Sequence: shared instrument, then boundary objects, then contact. Reversing the sequence produces the arithmetic in §20.7's second half.

SL-7 — Build a shared instrument. One measurement apparatus, readable by both sides without specialist tooling, owned by neither, current, and specified jointly. This is instrument 1 of 5 and first for a reason (§20.5).

SL-8 — Specify boundary objects, not forms. The artifact must tolerate different interpretations on each side. If it requires one side to translate into the other's frame before writing, it is a form; it reduces `C` but does not raise `D`.

SL-9 — Time-box the translator. Where a translator role is created, record its success condition as *reduction in `m` to the point where the role is unnecessary*, and record what Instruments 1 and 2 they are building. A permanent translator means the boundary cost was paid, not removed (§20.5).

SL-10 — Rotation is offered, never assigned. Declining shall carry no consequence to standing, assignment, or evaluation (CANON §5.4). **The safety property is the point of the method and an assigned rotation forfeits it.**

SL-11 — Rotation dosage is computed, not chosen. Use the §20.8 arithmetic to find the `rho` at which the step turns positive at your measured `m`, then convert to weeks (Worked Example). Under-dosing produces tourism; over-dosing at high `m` produces harm.

SL-12 — Co-locate on shared work, not adjacent work. Physical proximity without a shared output raises `C` and does not raise realised `u`. Verify by direct observation (Whyte's method, Chapter 12), not by seating plan.

SL-13 — Check Allport's four conditions explicitly before any contact-increasing intervention: equal status in the situation, common goals, cooperation not competition, institutional support. **Record which are absent.** Absent institutional support means `m` was not lowered (§20.9).

SL-14 — Prefer semilattice to tree. Where a person can belong to a function *and* a product *and* a guild, arrange for it. Overlapping membership realises `u_cross` without routing through management (Alexander, 1965; Chapter 8).

SL-15 — Do not remove the boundary by default. Boundaries are a `C`-control device and the quadratic is real. The question is whether the boundary is drawn where the seam should be and whether its crossing cost is proportionate (§20.6).

SL-16 — Nothing in this chapter is mandated. Every instrument shall be declinable, with the decline recorded and consequence-free. **A method whose worst outcome is a shrug must be kept that way, and any version that acquires a compliance component has left this manual's scope (CANON §5.4).**

SL-17 — No exhortation as intervention. Statements of value, collaboration campaigns, and culture messaging shall not be recorded as interventions, because they change no parameter (§20.3). They may accompany an intervention. They are not one.

SL-18 — Predict the reverse rate. For any cross-boundary change, state in advance how it will be detected if people drift back after attention moves on. **Forward rate alone cannot distinguish a lowered valley from raised agitation** (§20.3), and that distinction is the whole difference between a change and a churn.

SL-19 — Measure `D` by asking, not by asserting. The test of whether the cross-domain position is genuinely more comfortable is whether practitioners report it as such after using it. Comfort is read, not declared (Book One, Node 02).

SL-20 — Badge discipline. The landscape and escape-rate models are analogies and shall be labelled as such in any report. **An unbadged postulate is a defect** (Chapter 27).

The Failure Modes

F20.1 — THE EXHORTATION. *Clinical description.* Leadership declares collaboration a value. Nothing about the landscape changes. Crossing rate rises briefly and returns. *Parameter damaged.* None directly, which is why it persists — it is not visibly harmful. It damages future interventions by raising thresholds (Chapter 19, F19.3). *Diagnostic signature.* Forward crossing rate rises during the campaign and falls after; no instrument was built; the same campaign has run before under a different name. *Measurement.* SL-18: measure the reverse rate three months after attention moves on. *Fix.* Build something. §20.5, instruments 1 and 2.

F20.2 — FORCED INTEGRATION. *Clinical description.* A reorganisation raises `rho` sharply while `m` is untouched. Contact rises; cost rises faster; output falls. *Parameter damaged.* `C`, severely. Frequently drives `h_t` below zero, at which point Corollary 3.4 applies and **no interaction geometry can prevent decline.** *Diagnostic signature.* Meeting load up; delivery down; both sides

report the other side as the problem; time-to-decision rises; and — the tell — the harder people try, the worse it gets. *Measurement.* Estimate m (SL-4) and compare with m∗ (SL-5). If m > m∗, that is the finding and no further diagnosis is required. *Fix.* Reduce rho and spend the recovered budget on lowering m. **Reversing an integration is politically expensive and is the correct move.**

F20.3 — PROXIMITY WITHOUT CONTACT. *Clinical description.* Two domains seated together on adjacent, not shared, work. Neither u nor D moves. C rises from noise, interruption, and the loss of within-domain fluency. *Parameter damaged.* C up, realised u flat. *Diagnostic signature.* Cross-domain conversation counts near zero despite co-location; headphone use up; each domain has reconstructed a virtual boundary using calendars or channels. *Measurement.* Direct observation of actual cross-domain exchanges per day (Whyte's method). *Fix.* SL-12. Shared output, or separate them again and stop paying the acoustic cost.

F20.4 — THE PERMANENT TRANSLATOR. *Clinical description.* One person carries all cross-boundary meaning, indefinitely, competently. *Parameter damaged.* Robustness. m was paid, not lowered; the moment the translator leaves, rho collapses to its pre-translator value and the organisation discovers the boundary was never crossed. *Diagnostic signature.* All cross-boundary threads route through one name; that person is described as indispensable; no shared instrument exists. *Measurement.* Remove them for two weeks and count cross-boundary exchanges. *Fix.* SL-9. Give the role a termination condition and a build list.

F20.5 — THE DECLARED VALLEY. *Clinical description.* A cross-domain position is announced as better without being made better. D = 0. The barrier remains 1.0000. Movement is zero and is attributed to resistance. *Parameter damaged.* None, plus the credibility required for the next attempt. *Diagnostic signature.* Practitioners who tried it report it was slower, harder, or less understood; the initiative's materials describe benefits in the future tense. *Measurement.* SL-19. Ask the people who tried it. **Comfort is read, not declared.** *Fix.* Build the instrument that makes it actually faster, then stop talking about it.

F20.6 — THE TREE. *Clinical description.* A strict hierarchy in which every person belongs to exactly one category, so every cross-category path runs through management. *Parameter damaged.* u_cross realisation, structurally; and C, because the routed path is long. *Diagnostic signature.* Cross-domain questions escalate before they are answered; the answer arrives via two managers; nobody has a working relationship outside their branch. *Measurement.* Map the real working graph and count paths that traverse a common manager. *Fix.* SL-14. Overlapping membership — guilds, product teams, cross-cutting projects. **This is Alexander's semilattice and it is the structural form of building a lower valley.**

F20.7 — CONTACT WITHOUT ALLPORT. *Clinical description.* Contact increased with unequal status, no common goal, competitive framing, or no institutional support. The contact literature is cited in support. *Parameter damaged.* C up; u down where the higher-status frame absorbs the lower-status one. *Diagnostic signature.* One domain's vocabulary has become canonical; the other reports having "learned to speak their language" — which is u_cross being converted into u = 0 and reported as progress. *Measurement.* SL-13 checklist. Also test whether both sides' terms appear in shared artifacts, or only one side's. *Fix.* Restore the absent condition, especially institutional support, which is the one that lowers m (§20.9).

F20.8 — THE UNBUDGETED SEAM. *Clinical description.* Cross-boundary work is genuinely valued and has no budget line, no owner, and no time allocation, because budgets follow the tree. *Parameter damaged.* rho, absolutely, regardless of anyone's willingness. *Diagnostic signature.*

Everyone agrees the cross-boundary work matters; it is done in gaps by volunteers; it is the first thing dropped under load. *Measurement.* Look for the budget line; its absence is the finding. *Fix.* SL-1 and Conway. Create the line, or accept that the work will not happen and stop reporting it as a priority.

F20.9 — THE COERCIVE FIX. *Clinical description.* Cross-functional behaviour is added to performance evaluation. Compliance rises to near-universal. *Parameter damaged.* Everything the chapter is trying to protect. `u_cross` becomes performed rather than realised; the measurement detaches from the thing (Goodhart); and the method's safety property is destroyed, since declining now costs standing. *Diagnostic signature.* Cross-boundary artifact counts high; cross-boundary *decisions* unchanged; the artifacts are produced near review dates. *Measurement.* Compare artifact production against the evaluation calendar. *Fix.* Out of scope for repair by this manual (CANON §5.4). Remove the evaluation component and re-measure from a worse baseline than you started with.

Worked Example

The system

A **60-person product organisation** in two long-standing silos of 30: a **platform** group (infrastructure, tooling, reliability) and an **application** group (customer-facing features, domain workflows). Both are competent. Both are frustrated. Delivery has been slowing for two years.

Survey (Chapter 24 method).

```
N          = 6 distinct knowledge categories across the whole organisation
ln N       = ln 6 = 1.7918 nats
u_cross      = 0.55          between-silo non-parallelism, from a cross-
boundary disagreement audit
u_div    = 0.0990    within-silo dissent, same instrument applied inside each silo
u available= 0.6490     Theorem 4.1
delta      = 0.05       per month
Q_0     = 4.40     order proxy Y: composite of (i) fraction of interfaces with a
                  current written contract, (ii) mean time from cross-boundary
                      question to usable answer, inverted, (iii) rework rate
c_in       = 0.00012    cost per within-boundary pair in contact, normalised
within-boundary pairs = 2 x 435 = 870
cross-boundary pairs  = 30 x 30 = 900
```

Proxy justification (Paper E §6.3). Y is assembled from interface-contract currency, answer latency, and rework — all recorded from artifact and ticket histories. U comes from a disagreement audit on posed technical questions. **These share no input. A Y that is secretly U fits beautifully and means nothing.**

Measured `m` (SL-4). Twenty matched exchanges, ten within-boundary and ten cross-boundary, timed from question asked to answer usable. Median within: 1.8 hours. Median cross: 14.6 hours.

```
m ~ 8
```

Measured rho (SL-3). Cross-boundary pairs with any shared artifact, meeting, or working thread in the last quarter: 27 of 900.

```
rho = 0.03
```

Measured eta. Fraction of cross-boundary decisions traceable into a durable interface contract: `eta = 0.34`.

Case A — the status quo

```
u realised = 0.03(0.55) + 0.0990 = 0.1155
N_eff      = 1 + 5(0.1155)       = 1.5775
S          = ln(1.5775)          = 0.4558 nats
eta S      = 0.34 x 0.4558       = 0.1550

C = 0.00012(870) + 8(0.00012)(27) = 0.1044 + 0.0259 = 0.1303
delta Q = 0.05 x 4.40 = 0.2200

h    = 0.34(1.7918) − 0.2200 − 0.1303 = 0.6092 − 0.3503 = +0.2589
step = 0.1550 − 0.2200 − 0.1303 = −0.1953        DECLINING
Q_inf = (0.1550 − 0.1303)/0.05 = 0.4933
```

h > 0 and the step is negative — the same diagnostic split as Chapter 18's baseline. The regime is feasible; the geometry is not achieving it. Required surplus:

```
S_req = (0.2200 + 0.1303)/0.34 = 1.0304 nats  <  ln N = 1.7918
u_req = (e^{1.0304} − 1)/5 = (2.8021 − 1)/5 = 0.3604
```

And here is the finding that makes the whole diagnosis: u available on the roster is **0.6490**, comfortably above the required **0.3604**. **The organisation already owns enough difference to sustain itself. It is realising 0.1155 of 0.6490 — 17.8% — because 873 of its 900 cross-boundary pairs never speak.** Nobody needs to be hired. Nobody needs to be persuaded of anything. The difference is on the payroll and out of contact.

Left alone:

month	Q_t
0	4.4000
3	3.8428
6	3.3651
9	2.9555
12	2.6043
18	2.0451
24	1.6340

Two years of competent work by sixty competent people, and stored order falls by 63%.

Case B — forced integration (the intervention that gets chosen)

A reorganisation. Cross-functional squads, mandatory joint planning, shared standups. `rho` rises from 0.03 to 0.45 — a genuine, large, successfully executed increase in contact. **`m` is not touched**, because no shared instrument was built and no boundary object was specified: the two domains still measure different things and still mean different things by the same words.

```
u realised = 0.45(0.55) + 0.0990 = 0.3465     (up from 0.1155)
N_eff      = 2.7325                            (up from 1.5775)
S          = ln(2.7325) = 1.0052 nats          (up 121%)
eta S      = 0.34 x 1.0052 = 0.3418            (up 121%)

C = 0.1044 + 8(0.00012)(0.45 x 900) = 0.1044 + 0.3888 = 0.4932   (up 279%)

h    = 0.6092 − 0.2200 − 0.4932 = −0.1040        h < 0
step = 0.3418 − 0.2200 − 0.4932 = −0.3714       DECLINING FASTER THAN BEFORE
```

Surplus rose by 121% and the system got worse. The reorganisation did exactly what it set out to do — it more than doubled realised non-parallelism — and it was destructive anyway, because cost rose by 279% against a surplus that rises logarithmically.

And h has crossed zero. **Corollary 3.4 now applies in its hardest form: `C_t > eta ln N − delta Q_t`, therefore no choice of interaction geometry can prevent order from decreasing.** At this point there is no seating plan, no meeting protocol, no facilitation, and no amount of goodwill that recovers the system. The only available moves are to cut `C`, raise `eta`, or shed `Q`.

Check against `m*`: at `rho = 0.45` and `eta = 0.34`, `m* = 3.17`. The organisation is operating at `m = 8`. **It was above the critical multiplier before the reorganisation began, and the reorganisation's entire mechanism was to increase the quantity that `m` multiplies.**

month	`Q_t` under forced integration
0	4.4000
2	3.6757
4	3.0221
6	2.4321
8	1.8997
10	1.4192
12	0.9856

Twelve months, 78% of stored order gone, against 24 months and 63% for doing nothing. **The intervention is worse than inaction by a wide margin, and everyone involved is working harder.**

Case C — build the lower valley

The same target `rho = 0.45`, approached in the opposite order (SL-6): **lower `m` first.**

Three things are built, over four months, before any contact target is set:

1. **A shared instrument** (SL-7): one service-level dashboard, one definition of latency, error, and saturation, one dataset, owned by neither group, readable by both without specialist tooling. Most cross-boundary arguments in the sampled history were arguments about what was happening. Those now terminate at a shared screen.
2. **A boundary object** (SL-8): a written interface contract per integration point, plastic enough that platform reads it as an operational commitment and application reads it as a behavioural guarantee, without either translating into the other's frame.
3. **Two rotations, offered not assigned** (SL-10): four people, one quarter each, working — not observing — inside the other domain. Six people declined and the declines were recorded with no consequence, per SL-16.

Measured after four months: `m` falls from 8 to **2** (median cross-boundary answer latency 14.6 hours to 3.6 hours). `eta` rises from 0.34 to **0.46**, because the interface contract is a durable artifact and cross-boundary decisions now land in it.

```
u realised = 0.45(0.55) + 0.0990 = 0.3465
S          = ln(2.7325) = 1.0052 nats           (identical to Case B)
eta S      = 0.46 x 1.0052 = 0.4624              (Case B: 0.3418)

C = 0.1044 + 2(0.00012)(405) = 0.1044 + 0.0972 = 0.2016    (Case B: 0.4932)

h    = 0.46(1.7918) - 0.2200 - 0.2016 = 0.8242 - 0.4216 = +0.4026
step = 0.4624 - 0.2200 - 0.2016 = +0.0408        GROWING
Q_inf = (0.4624 - 0.2016)/0.05 = 5.2160

S_req = (0.2200 + 0.2016)/0.46 = 0.9165  ->  u_req = 0.3001
realised u = 0.3465 >= 0.3001                      FEASIBLE, with margin
```

Case B and Case C have identical contact, identical realised u, and identical surplus. The only differences are `m` and `eta`. Case B declines at -0.3714 per month; Case C grows at +0.0408. The entire difference between a destructive reorganisation and a working one is the boundary cost and the conversion rate — neither of which is about how much people collaborate.

month	Q_t (Case C)	h_t	step
0	4.4000	+0.4026	+0.0408
3	4.5164	+0.3968	+0.0350
6	4.6162	+0.3918	+0.0300
9	4.7017	+0.3875	+0.0257
12	4.7751	+0.3839	+0.0220
18	4.8919	+0.3780	+0.0162
24	4.9777	+0.3737	+0.0119

At month 24: **1.6340** doing nothing, **under 0.99 at month 12** under forced integration, **4.9777** under the lower-valley method, converging on 5.2160. And the step shrinks as `Q` rises — success consuming its own headroom, Banach's contraction arriving at Theorem 3.5's fixed point.

The dosage calculation

`rho = 0.45` was not chosen. It was computed. At `m = 2` and `eta = 0.46`, sweeping `rho`:

rho	C	u	S	step at Q = 4.40
0.00	0.1044	0.0990	0.4021	-0.1394
0.05	0.1152	0.1265	0.4901	-0.1097
0.10	0.1260	0.1540	0.5710	-0.0833
0.20	0.1476	0.2090	0.7154	-0.0385
0.30	0.1692	0.2640	0.8416	-0.0021
0.45	0.2016	0.3465	1.0052	**+0.0408**
0.60	0.2340	0.4290	1.1458	+0.0731
0.80	0.2772	0.5390	1.3070	+0.1040
1.00	0.3204	0.6490	1.4457	+0.1246

The step turns positive just above `rho = 0.30`. `rho = 0.45` is chosen as the first target because it clears the crossing with margin, and higher targets are held in reserve — note the marginal gain from 0.45 to 1.00 is `0.1246 − 0.0408 = 0.0838` for more than double the contact, which is surplus concavity showing up as a diminishing return on integration.

The same sweep at the un-remediated `m = 8`, `eta = 0.34`:

rho	C	S	step at Q = 4.40
0.00	0.1044	0.4021	-0.1877
0.10	0.1908	0.5710	-0.2167
0.20	0.2772	0.7154	-0.2540
0.30	0.3636	0.8416	-0.2975
0.45	0.4932	1.0052	-0.3714
0.60	0.6228	1.1458	-0.4532
0.80	0.7956	1.3070	-0.5712
1.00	0.9684	1.4457	-0.6968

Monotonically worse at every level of contact, all the way to complete integration. There is no value of `rho` at which the un-remediated organisation benefits from more collaboration. That is what `m > m*` means, and it is why SL-5 forbids raising contact before measuring it.

Rotation dosage in weeks

Converting `rho` into a schedule, at `m = 2` and `eta = 0.46`, `Q = 4.40`:

weeks per person per year across the boundary	rho	u	S	C	step
0	0.0000	0.0990	0.4021	0.1044	-0.1394

weeks per person per year across the boundary	rho	u	S	C	step
1	0.0192	0.1096	0.4369	0.1086	-0.1276
2	0.0385	0.1202	0.4705	0.1127	-0.1163
4	0.0769	0.1413	0.5345	0.1210	-0.0952
8	0.1538	0.1836	0.6513	0.1376	-0.0580
13	0.2500	0.2365	0.7805	0.1584	-0.0194
26	0.5000	0.3740	1.0543	0.2124	**+0.0526**

A one-week or two-week rotation is tourism: the step improves by 0.012 to 0.023 and remains firmly negative. Even a full quarter — thirteen weeks, which most organisations would consider a serious commitment — leaves the step at -0.0194. **Rotation alone, at any dose an organisation will tolerate, does not save this system.** It is a supporting instrument. The shared instrument and the interface contract are what move `m` and `eta`, and they are what make the modest `rho` affordable.

This is the practical reason instrument 1 outranks instrument 4 in §20.5, and it is the kind of conclusion that only appears once the arithmetic is done.

Falsification statement, pre-registered

```
MODEL      Y_{t+1} = a Y_t + b F(U_t) − Chat_t + e_t ,  0 <= a <= 1,  b >= 0
FIT        on months −18..0, then FROZEN
PREDICT    months 1..24 held out

PRE-REGISTERED PREDICTIONS
 P1  median cross-boundary answer latency falls below 5.0 hours within 4 months
     of the shared instrument going live (m from ~8 to <=3).
 P2  h_t stays positive throughout; Y rises monotonically from month 1.
 P3  Y_24 lies in [4.7, 5.2]; Y_infinity estimate in [4.9, 5.5].
 P4  the rotation declines (SL-10) do not correlate with subsequent evaluation
     outcomes. This is an ETHICS check, not a performance one, and it is the
     single most important line in this preregistration.

BASELINES  (i) persistence  (ii) direct-input, raw U with no log
           (iii) matched autoregressive, no interaction term
           (iv) SHUFFLED-INPUT CONTROL

FALSIFIED IF
 1. the shuffled-input control predicts held-out Y as well as the ordered series;
 2. the matched autoregressive baseline is not beaten;
 3. fitted a outside [0,1] or fitted b negative;
 4. m does not fall below 3 despite the shared instrument being in daily use
    — which would mean the boundary cost is not where this chapter located it.
```

Falsifier 4 is the one worth stating loudest. The entire chapter rests on the claim that the boundary cost is dominated by translation and disputed measurement, and that a shared instrument removes most of it. If `m` stays at 8 while a well-built, well-used shared instrument sits on the wall, then the boundary cost lives somewhere this chapter did not look — in incentives, in status, in scheduling, in something else — and the recommendation to build instruments first is wrong for that system. **A framework that cannot fail cannot be trusted, and this is exactly where this one fails.**

And the final observation, which is the chapter's reason for preferring this method over every alternative: **in Case C, six people declined the rotation and nothing happened to them.** The dashboard was built and they could ignore it. The interface contracts existed and they could read them or not. The system improved from 4.4000 to 4.9777 with six recorded refusals in it. **A method that works by removing resistance produces a system whose improvement is compatible with people saying no — and that is not a concession the method makes. It is the property that makes it safe to recommend.**

Chapter 21 . Order Without a Centre

The Law

There is a habit of mind so old that it is rarely stated as a proposition, and therefore rarely tested: **that order requires an orderer.** Where there is a pattern, look for the patterner. Where there is a schedule, find the scheduler. Where a hundred people arrive at a coherent result, locate the person who held the plan.

The habit is not stupid. It is often correct. But it is not a law, and this chapter is about the enormous class of systems in which it is false — and about how to build them on purpose.

21.0.1 What is proven

Everything in Part I applies unchanged. The storage law does not know whether a coordinator exists:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t        S_t = F(U_t) = ln(1 + (N−1) U_t)
```

and the master gauge is unchanged:

```
h_t = eta ln N − delta Q_t − C_t
```

What changes when you remove the centre is not the equation. It is the **parameters**. A coordinator is not outside the system; a coordinator is a particular setting of `(delta, eta, C)` and a particular *location* for `Q`. That is the entire analytical content of this chapter, and it is worth stating in its strongest form before any biology or sociology is introduced:

> **A coordinator is a device that raises eta at the price of C, and concentrates Q in a single store whose delta is the coordinator's own turnover rate.**

Each clause is checkable. A coordinator raises `eta` because someone is paid to convert loose interaction into durable record — that is what the role *is*. A coordinator raises `C` because the role costs money, attention, scheduling, and queueing. And a coordinator concentrates `Q`, because the state of the system now lives substantially inside one person's working memory, calendar, and relationships. When that person leaves, the system does not lose a worker. It loses a fraction of its stored order in a single step.

21.0.2 A derived identity: headroom at equilibrium is unused variety

One algebraic consequence of Paper E is not stated in the paper and is used repeatedly below, so it is derived here in full. It is a consequence, not a new theorem.

At any steady state of the storage law, `Q_{t+1} = Q_t`, so from the recursion:

```
Q = (1 − delta) Q + eta S − C
0 = −delta Q + eta S − C
delta Q + C = eta S
```

Substituting into the headroom gauge:

```
h_inf = eta ln N − delta Q − C
      = eta ln N − eta S
      = eta ( ln N − ln N_eff )
```

```
= eta ln ( N / N_eff )
```

At equilibrium, headroom equals the conversion efficiency times the log of the redundancy still present in the roster. Headroom in a settled system is exactly the surplus you are not collecting. It is not slack in the budget; it is the measure of the difference you have and are failing to use.

Two corollaries follow immediately and are used in §21.6 and in Chapter 23:

1. A system at equilibrium with `u = 1` — perfectly orthogonal participants — has `h_inf = 0`. It is running at its ceiling. It is stable and it cannot grow. Nothing is wrong with it.
2. A system at equilibrium never has negative headroom. **Negative headroom is always a disequilibrium condition**: it means the stored order was accumulated under one parameter set and is now being carried under another. Chapter 23 is entirely about that case.

This identity holds regardless of how `C` depends on `Q`, because it only used the definition of a steady state. It is verified numerically twice in the Worked Example below.

21.0.3 What is postulated

The following are **design postulates**, badged here and every time they are leaned on. They are not derived from Papers A–E and they are not empirical findings. They are engineering heuristics that the law makes *legible*, which is a weaker and more honest claim.

- **Postulate M1 (state-in-medium).** *Postulate:* moving the system's state out of a coordinator and into a shared, externally visible medium lowers `delta` and lowers `C` more than it lowers `eta`. Net effect: headroom rises.
- **Postulate M2 (medium quality).** *Postulate:* the `eta` of a stigmergic system is dominated by five properties of its medium — visibility, writability, durability, low friction, and append-orientation — and these can be specified and inspected directly.
- **Postulate M3 (centre as hub).** *Postulate:* a coordinating centre is a topological hub, and therefore inherits the robustness profile Albert, Jeong & Barabási (*Nature* 2000) measured for hubs in scale-free networks: robust to random failure, fragile to targeted removal.

M1 and M2 are the ones that carry the chapter. They can be falsified in any real installation by measuring `delta`, `eta` and `C` before and after, by the protocol in Chapter 24. **If you install a shared medium and headroom does not rise, the postulate failed in your system and you should say so.**

The Engineering

21.1 Stigmergy: Grassé, 1959

Pierre-Paul Grassé, studying nest construction in termites of the genus *Bellicositermes*, faced a problem that his field had been quietly avoiding. A termite mound is a coherent structure with ventilation shafts, brood chambers and a royal cell. It is built by hundreds of thousands of insects, none of which possesses a representation of the mound. There is a queen, but the queen is an ovary, not an architect. No termite has seen the finished thing. No termite is told what to do.

Grassé's answer, published in 1959, was that the workers do not coordinate with each other at all. They coordinate with **the work already done**. A termite deposits a pellet of soil impregnated with pheromone. That deposit changes the local environment. The changed environment is the stimulus that makes the next deposit more likely at that spot, which changes the environment further. The structure grows because **the state of the project is stored in the project**, and each worker reads the state by looking at it.

He named it *stigmergy*, from Greek *stigma* (mark, sign) and *ergon* (work): the work is the sign.

The definition to hold on to is this:

> **Stigmergy: coordination in which agents communicate only by modifying a shared medium, and no agent holds the plan.**

Note what stigmergy is not. It is not "self-organisation" as a mood. It is not emergence as mystique. It is a specific mechanism with two named requirements — a medium that agents can both read and write, and a local rule that maps what is read to what is written. Both requirements are engineerable. Both are frequently absent in organisations that claim to be decentralised.

The ant trail is the same mechanism in one dimension. An ant returning from food lays pheromone. Pheromone attracts. More traffic means more pheromone; pheromone evaporates, so a trail to an exhausted source fades. The trail network is not stored in any ant. It is stored on the ground, and its `delta` is literally the evaporation rate of the trail chemical.

That last point deserves emphasis because it is the whole book in miniature. **The evaporation of the pheromone is not a defect of the ant trail. It is the mechanism by which the colony forgets a food source that is gone.** A pheromone that did not evaporate would build a permanent motorway to an empty patch of ground. This is Chapter 3's argument arriving from entomology: `delta = 0` is not perfect memory, it is the inability to learn.

Two honest limits on the biology, stated now so they are not smuggled in later:

- Termite and ant behaviour is used here as an **existence proof of a mechanism**, not as a model of human organisations. Humans read media in ways ants do not: they infer intention, they anticipate, they lie, and they defect strategically. Every human stigmergic system therefore needs the governance layer of §21.6, which the termites do not need.
- Grassé described stigmergy; he did not quantify it in the parameters of Papers A–E, which did not exist. The mapping from stigmergy to `(delta, eta, N, u, C)` in this chapter is **this manual's modelling choice**, badged as such.

21.2 Two human media that work

The version-controlled repository. A distributed software project of a thousand contributors has no one who has read all the code. The state of the project is the repository: the current tree, the history, the open issues, the test results. A contributor reads the state by cloning it, modifies it, and offers the modification back. The medium is visible, writable via a defined proposal mechanism, durable, low-friction, and append-oriented. Maintainers exist, but they are not coordinators in the sense that matters: they do not hold the plan. If every maintainer vanished on Tuesday, the state of the project would be entirely intact on Wednesday — which is precisely what is *not* true of a project whose state lives in a programme manager's spreadsheet.

The encyclopedia anyone can edit. An article's state is the article, plus its history, plus its talk

page. No editor has read the encyclopedia. There is no assignment of topics. The coordination signal is the visible deficiency of the current text — a missing citation, a contradiction, an obvious gap — which functions exactly as Grassé's pheromone-laden pellet: the work already done is the stimulus for the next increment.

Notice that in both cases the medium carries a **grammar of incompleteness**. A repository has failing tests and open issues; an article has "citation needed" and a stub tag. These are not administrative overhead. **A stigmergic medium must be able to display its own unfinished edges, or there is nothing for a contributor to respond to.** A medium that only shows finished states is not a stigmergic medium; it is a report.

Both examples come with a warning that §21.9 develops: neither is centreless in the naive sense. Both have governance, and both spent years acquiring it, usually after a crisis. A shared medium is necessary and it is not sufficient.

21.3 Schelling: coordination without any communication at all

Thomas Schelling's *The Strategy of Conflict* (1960; Nobel 2005) isolates the limiting case. Two people must meet in a city on a given day. They cannot communicate. Where do they go?

Empirically, people do far better than chance. They converge on **focal points** — solutions that are salient, and that each person expects the other to find salient. The grand central landmark. Noon. The round number. The obvious default.

A focal point is a stigmergic medium with zero write bandwidth. Nobody modified anything; the coordination is carried entirely by shared, publicly evident structure. **Before you build any medium at all, ask what the participants already share that is salient, because salience is free coordination.** Focal-point engineering is the cheapest intervention in this chapter:

Focal device	What it removes
A fixed weekly time that never moves	The scheduling negotiation, permanently
A single canonical location for a class of document	The "where does this live" exchange
A default that is correct 80% of the time	Eighty per cent of the decisions
A round-number threshold ("anything over 10k")	The case-by-case escalation judgement
A naming convention	The lookup

Each of these is a `C` reduction, and `C` reductions convert directly into headroom by the gauge. They cost approximately nothing and they are almost always available in a system that has never looked for them. Schelling's contribution to this manual is the observation that **the most efficient coordination is the coordination you never have to perform.**

21.4 What makes a good shared medium

This is the specification-grade content of the chapter, and it is where Postulate M2 does its work. Five properties, each with an inspection test that a surveyor can run in an afternoon.

V — Visible Any participant can see the current state without asking permission and without asking a person.

Inspection test: pick three participants at random, at three different levels of the system. Ask each to show you the current state of one live item, right now, from where they are sitting. Time it. **If the median time exceeds sixty seconds, or if any of the three has to ask someone, the medium fails V.**

Visibility is the property most often lost by accident. A medium starts open, then acquires access controls for good reasons, then acquires so many that reading it requires knowing whom to ask. At that moment you no longer have a shared medium; you have a coordinator wearing a database costume.

W — Writable Any participant who observes something can record it, and the recording takes effect without adjudication.

Inspection test: count the steps between "a participant notices something" and "the medium shows it." **Two steps or fewer passes. Four or more fails.** Count review as a step only if it is blocking; asynchronous review after the fact is not a write barrier.

Writability is where most organisational media die. Read-only dashboards are extremely common and they are not stigmergic media at all: they are broadcast. Broadcast has `eta` close to zero for the receivers, because nothing a receiver learns can be deposited back.

The subtlety: writability does not mean unmoderated. It means the *default* is write-through and the correction is after the fact. An append-only log with revert is writable. A submission queue with a two-week approval cycle is not.

D — Durable The state survives the departure of any individual, and survives the interval between uses.

Inspection test: the **bus-count**. For each critical piece of state, count how many people would have to leave simultaneously for it to become unrecoverable. **A bus-count of 1 for any critical item is a failure.** Then run the **fortnight test**: is state recorded two weeks ago still intelligible today to someone who was not present when it was written?

Durability is the direct `delta` control. This is the parameter that a coordinator-centred system gets catastrophically wrong, because a person's memory has a `delta` that looks small in the short run and equals 1.0 on the day they resign.

L — Low friction The cost of one read and one write is small relative to the value of the increment.

Inspection test: measure the seconds to complete one typical write, end to end, including authentication, navigation, and any required fields. **Under thirty seconds passes. Over two minutes fails, and will be routed around.** Then count required fields: every mandatory field is a tax on every write, forever.

Friction is not a minor UX concern. It is the term that decides whether the medium is used at all, and an unused medium has `eta = 0` no matter how well designed. The characteristic failure is

the elaborate system that captures everything perfectly and is filled in by nobody, sitting beside a group chat that captures nothing and holds the real state.

A — Append-oriented New information is added; old information is superseded rather than erased; the history is retrievable.

Inspection test: can you reconstruct the state as of a named date last quarter? **If not, the medium is not append-oriented**, and every dispute about what was decided will be resolved by whoever speaks most confidently.

Append-orientation is what makes the medium adjudicative. It is also what makes deliberate forgetting possible *as a decision* rather than as an accident — you can retire a rule explicitly, and the retirement is itself a record. Chapter 23 depends on this.

The mnemonic is **VWDLA**, and the survey form in the Specification scores each from 0 to 4.

21.5 Ostrom's eight design principles, in full

Elinor Ostrom's *Governing the Commons* (Cambridge University Press, 1990) is the most directly engineerable document in the sociology of shared resources, and it earned her the 2009 Nobel Memorial Prize in Economic Sciences. Its method matters as much as its result: rather than theorising about whether commons must fail, Ostrom and her collaborators assembled and compared a large body of field cases of common-pool resources — Swiss alpine pastures, Japanese village forests, Spanish *huerta* irrigation systems, Philippine irrigation federations, inshore fisheries, groundwater basins — and asked what distinguished the arrangements that persisted for centuries from those that collapsed.

A **common-pool resource** in her sense has two features: it is *subtractable* (what one person takes, another cannot), and it is *hard to exclude people from*. Grazing land, fish stocks, groundwater, irrigation water, a shared workshop, a shared codebase, a shared meeting-hour budget, the attention of a small team.

The eight principles she identified as characterising long-enduring, self-governed CPR institutions are given below in full. **Ostrom did not claim these are sufficient, and she did not claim they are a recipe.** She characterised them as features found in the enduring cases and absent or weak in the failed ones. That is the claim this manual repeats, and it is not a stronger one.

Principle 1 — Clearly defined boundaries. Individuals or households with rights to withdraw resource units from the common-pool resource must be clearly defined, as must the boundaries of the resource itself.

Why it matters: if you cannot say who is in and what the resource is, you cannot say what over-use looks like, and everyone's estimate of everyone else's restraint is a guess.

Parameter: `N` and `C`. Undefined boundaries make the effective category count unknowable and generate continuous boundary-dispute cost.

Principle 2 — Congruence between appropriation and provision rules and local conditions. Appropriation rules restricting time, place, technology, and/or quantity of resource units are related to local conditions and to provision rules requiring labour, material, and/or money.

Why it matters: two congruences are required, and they are usually confused. First, the rules must fit the *physical* conditions — a rule written for a wet valley will destroy a dry one. Second, what you may *take* must be proportionate to what you *put in*. A system where extraction rights are detached from contribution obligations is a system that has instructed its members to free-ride.

Parameter: `C` and `eta`. Ill-fitting rules generate constant exception-handling; contribution proportionality is what keeps the input term funded.

Principle 3 — Collective-choice arrangements. Most individuals affected by the operational rules can participate in modifying the operational rules.

Why it matters: this is the principle most often violated by well-intentioned outsiders, who write good rules for people who had no hand in them. Rules written by the users are better fitted, because the users hold the local knowledge; and they are better complied with, because compliance with a rule you helped write is not obedience.

Parameter: `u` and `eta`. Participation in rule-making is exactly the mechanism that puts multiple distinct viewpoints into the rule, raising `u`; and a rule you helped write is a rule that has already converted into your own practice, which is `eta`.

Principle 4 — Monitoring. Monitors, who actively audit common-pool-resource conditions and appropriator behaviour, are accountable to the appropriators or are the appropriators themselves.

Why it matters: this is Ostrom's sharpest empirical finding and the one that overturns the standard framing. The enduring commons are not populated by unusually virtuous people. They are populated by ordinary people who **can see what everyone else is doing.** Monitoring is not surveillance imposed from outside; in the durable cases it is mutual, cheap, and largely a by-product of the work itself — the irrigator who opens the gate can see the level in the channel, and can see who opened it last.

Parameter: `C` directly, via the free-riding loss term. Monitoring is a cost that pays for itself by removing a larger cost. The Worked Example computes this trade explicitly.

Principle 5 — Graduated sanctions. Appropriators who violate operational rules are likely to be assessed graduated sanctions — depending on the seriousness and context of the offence — by other appropriators, by officials accountable to these appropriators, or by both.

Why it matters: graduation is the load-bearing word. A system with one severe sanction has, in practice, no sanction: the penalty is so disproportionate to a first small lapse that nobody applies it, everybody knows nobody applies it, and the rule becomes advisory. A first offence that draws a word, a second that draws a small fine, a third that draws exclusion, is a system in which the *first* step is actually taken. **Graduated sanctions are cheap precisely because they are mild, and effective precisely because they are used.**

Parameter: `C`. Graduation minimises the cost of enforcement per unit of compliance obtained, and it prevents enforcement itself from becoming the system's dominant conflict source.

Principle 6 — Conflict-resolution mechanisms. Appropriators and their officials have rapid access to low-cost local arenas to resolve conflicts among appropriators or between appropriators and officials.

Why it matters: every rule is ambiguous at some boundary, and every commons generates disputes. The question is never whether disputes occur but how expensive they are. *Rapid* and *low-cost* and *local* are all three doing work: a dispute mechanism that takes six months, costs a lawyer, or requires travelling to a capital is a mechanism that will not be used, and unresolved disputes compound.

Parameter: `C`. Unresolved conflict is a permanent, growing dissipation term, and it is the one most likely to be omitted from a cost tally because it is nobody's line item.

Principle 7 — Minimal recognition of rights to organise. The rights of appropriators to devise their own institutions are not challenged by external governmental authorities.

Why it matters: self-governance that can be overridden at any moment is not self-governance; it is a delegation that will be revoked at the first inconvenience. The requirement is *minimal* recognition — Ostrom's cases did not need the external authority to endorse or fund the local institution, only to refrain from voiding it.

Parameter: `eta`. If local decisions can be nullified from outside, interaction does not convert into durable structure. This is `eta` collapse by decree, and it is the single most common way a head office destroys a functioning local commons while believing it is helping.

Principle 8 — Nested enterprises. For common-pool resources that are parts of larger systems: appropriation, provision, monitoring, enforcement, conflict resolution, and governance activities are organised in multiple layers of nested enterprises.

Why it matters: it is the answer to scale. A single governance body for ten thousand people cannot monitor, cannot resolve disputes locally, and cannot fit rules to local conditions. Nesting lets each layer do what it can do at its own scale, with the layer above handling only what the layer below genuinely cannot.

Parameter: `C`, decisively, and it is computable. Brooks's communication-path count `n(n−1)/2` grows quadratically while surplus grows as `ln N`. With per-path cost 0.02 and `eta = 0.5`, headroom at `Q = 0`:

n	paths	C	eta.ln N	h
4	6	0.120	0.6931	+0.5731
5	10	0.200	0.8047	**+0.6047**
6	15	0.300	0.8959	+0.5959
8	28	0.560	1.0397	+0.4797
10	45	0.900	1.1513	+0.2513

n	paths	C	eta.ln N	h
11	55	1.100	1.1989	+0.0989
12	66	1.320	1.2425	**-0.0775**
16	120	2.400	1.3863	-1.0137
24	276	5.520	1.5890	-3.9310

Headroom peaks at n = 5 and goes negative between 11 and 12. A single flat body of 24 has h = −3.93: it cannot build anything, ever, at any interaction geometry.

Now nest it. Four cells of six, each with h = +0.5959, plus a federation layer of four delegates with h = +0.5731. *Modelling assumption, badged:* the law is stated per system, and summing headroom across cells assumes the cells are genuinely independent in their cost structure — if they are not, the sum overstates. Under that assumption the nested form has roughly 4 x 0.5959 + 0.5731 = 2.957 against the flat form's −3.931.

Principle 8 is not an administrative preference. It is the difference between a positive and a negative headroom budget, and the arithmetic is elementary.

21.5.1 The install order, and a correction to naive practice

The eight principles are a set, not a sequence, and the temptation on reading them is to start with the one that feels most like community: build trust, and the rest will follow. **Chapter 49 computes that this is the wrong order, and in a system that is already losing it is actively harmful.**

The mechanism is Theorem 3.5 and nothing more. Q_inf = (eta F(u∗) − c∗) / delta. Trust acts on delta — a group that relies on unverified commitments loses less order between periods. **But when the numerator is negative, dividing by a smaller positive denominator makes the result more negative.** In Chapter 49's worked consortium the three moves were separated and computed:

move	Q_inf
baseline	-0.0836
trust alone (delta 0.4300 -> 0.1700)	**-0.2114 — worse**
reciprocity alone (eta 0.1721 -> 0.4026)	+0.7394
both	+1.8701

Trust does not add order. It amplifies whichever sign the system already has. Halving delta on a negative numerator multiplied the loss by the same factor it would have multiplied a gain.

The design rule that follows is exact: **fix the numerator first — raise eta through reciprocity, or cut C — until eta F(u∗) > c∗, and only then invest in trust.** Read against Ostrom's list, this puts P4 (mutual monitoring) and P5 (graduated sanctions) ahead of anything that works by extending credit, because §21.6 shows both acting on the numerator — P4 on eta and on the free-riding term in C, P5 on C — and it puts P8 (nesting) ahead too, because §21.5 shows it changing the exponent on C. Chapter 49 §49.14 notes the failure signature to watch for: a trust programme in a body that converts nothing durable produces **a slower, more polite decline, and survey**

results saying the intervention worked. Cross-reference Chapter 49 §49.14 and Chapter 26's "stabilise before growing".

21.6 The two named failure modes of the commons, and Ostrom's answers

The centreless case has two classical objections. Both are real. Both have answers in the eight principles, and the answers are structural rather than moral.

The tragedy of the commons The argument in its standard form: a shared pasture, herders each adding animals. The benefit of an extra animal accrues entirely to its owner; the cost of the extra grazing is spread across all herders. Each herder therefore faces a private incentive to add animals until the pasture is destroyed. Rational individual action, collective ruin.

The arithmetic is correct for the situation described. **What the argument actually models is not a commons but an unmanaged open-access resource** — a resource with no boundaries, no rules, no monitoring, no sanctions and no communication. Ostrom's field cases are what happens when those things exist, and her central empirical point is that they very frequently do: communities across centuries and continents have built exactly the institutions the model assumes away.

The structural answers, by principle:

Mechanism of tragedy	Principle that addresses it
Anyone may take, so no restraint is credible	**P1** boundaries: define who may withdraw and from what
No limit on how much may be taken	**P2** congruence: appropriation limited and tied to provision
Restraint is invisible, so it is unrewarded	**P4** monitoring: restraint becomes observable
Over-use carries no consequence	**P5** graduated sanctions
Rules imposed from outside do not fit the pasture	**P3** collective choice, **P7** right to organise
Disputes over the boundary destroy the arrangement	**P6** low-cost local arenas
The pasture is bigger than any one village	**P8** nesting

The point to carry: **the tragedy is a property of a missing institution, not of human nature.** Where the eight features are present, the tragedy is not observed; where they are absent, it is. That is a claim about the world, from field data, and it is the claim Ostrom's prize was for.

Free riding The narrower and more common problem: the resource is not destroyed, but a fraction of participants consume the shared order without contributing to it. Nothing collapses. Everything just costs more, degrades faster, and quietly demoralises the contributors.

In the storage law, free riding is unambiguously a `C` term. Every unit of order consumed without replacement is a dissipation. Model it as `C = C_base + lambda * phi`, where `phi` is the free-riding fraction and `lambda` the loss per unit of it.

Take the tool library of the Worked Example: `eta = 0.55`, `delta = 0.12`, `N = 8`, `u = 0.62`, `S = 1.6752`, `C_base = 0.40`, `lambda = 0.9`.

Without monitoring, suppose `phi = 0.35`:

```
C = 0.40 + 0.9(0.35) = 0.7150
h(Q=0) = 0.55(2.0794) − 0.7150 = +0.4287
Q_inf = (0.55 x 1.6752 − 0.7150) / 0.12 = 1.7198
```

With Ostrom P4 and P5 — mutual monitoring at a cost of 0.08, graduated sanctions bringing `phi` down to 0.06:

```
C = 0.40 + 0.08 + 0.9(0.06) = 0.5340
h(Q=0) = 0.55(2.0794) − 0.5340 = +0.6097
Q_inf = (0.9214 − 0.5340) / 0.12 = 3.2281
```

Monitoring costs 0.08 and returns 0.261 in reduced free-riding loss. **Steady-state stored order nearly doubles — 1.72 to 3.23 — for a monitoring cost of eight hundredths of a unit.** This is why Principle 4 is the one to install first if you can only install one: it has the best ratio in the list, and it is the one people are most squeamish about because it sounds like surveillance.

It is not surveillance when it satisfies Ostrom's actual condition: **the monitors are accountable to the appropriators, or are the appropriators.** Mutual visibility among peers who share the resource is a different object from observation by an authority that does not. The first raises `eta`; the second raises `C` and lowers `u`, because people who are watched by an authority converge on the behaviour the authority rewards.

21.7 Where central enforcement is absent: Gambetta, and Putnam

Two field observations complete the picture, and both are uncomfortable in ways that make them useful.

Gambetta (1993) studied Sicily and asked what fills the gap when central enforcement of contracts is unreliable. His answer was that a market in private protection emerges: a networked structure of obligation, reputation and guarantee that does, in a fashion, produce order. Contracts are honoured. Disputes are settled. Coordination occurs.

The lesson for this chapter is deliberately double-edged. **Order without a centre is not automatically good order.** The mechanism this manual recommends — put the state in the medium, let obligation and reputation do the enforcement — is the same mechanism, structurally, that produces outcomes nobody would design. The difference is not in the topology. It is entirely in whether the governance features of §21.5 are present, and specifically P3 (can those affected change the rules?), P5 (are sanctions graduated, or terminal?) and P6 (is there a low-cost arena, or only force?). A centreless system that fails P3, P5 and P6 is not a commons. It is a protection racket with good uptime.

Carry this as a standing check. **Centrelessness is a topology, not a virtue.**

Putnam (1993), in *Making Democracy Work*, examined the divergent performance of regional governments in Italy after identical institutional reforms were applied across the country. The same formal institutions produced markedly different results north and south, and the difference tracked what he called social capital — dense horizontal associational life, choral societies, football clubs, mutual-aid associations, and the norms of reciprocity and generalised trust that go with them.

In this manual's terms: **identical formal structure, different `Q`.** The reform supplied a constitution; it did not supply stored order, and stored order cannot be legislated into existence in one step because the storage law integrates. `Q_t = (1−delta)^t Q_0 + sum of driven terms`. If `Q_0` is low, the transient is low, and the driven term needs time at a positive step to accumulate. Putnam's contested claim about the historical depth of the divergence is exactly a claim about how long that integration has been running.

Honesty note. Putnam's causal direction has been disputed at length — whether civic association produces good government or good government produces civic association is not settled by the data he presents, and this manual does not need it settled. What is used is the narrow and well-supported observation that **formal rules and stored social order are different quantities, and installing the first does not install the second.**

21.8 The topology argument

Albert, Jeong & Barabási (*Nature*, 2000) measured the robustness of scale-free networks and found an asymmetry: such networks tolerate random node failure remarkably well, and fail catastrophically under targeted removal of high-degree hubs.

Under Postulate M3, a coordinator is a hub — the highest-degree node in the coordination graph, by construction. The consequence is not that coordinators are bad. It is that a coordinated system has a **specific and predictable fragility profile**: it will handle random absences well and the targeted loss of the coordinator very badly, and this is a structural property that no amount of diligence by the coordinator can remove. Indeed diligence makes it worse, because the more effectively the coordinator holds the state, the more of it is lost when they go.

Two engineering consequences:

1. **Never treat coordinator loss as a personnel risk.** It is an architectural one. The mitigation is not a better succession plan; it is moving state out of the person and into the medium, so that the hub carries traffic and not storage.
2. **Measure hub concentration directly.** For each critical class of state, record the number of people who could reconstruct it. The distribution of that number across all state is the system's real robustness curve, and a system with many items at bus-count 1 is fragile no matter how good its documentation policy claims to be.

21.9 The honest limits of centrelessness

Four things a shared medium does not do, each stated because the literature on decentralisation routinely omits them.

It does not remove governance; it relocates it. The repository has maintainers and a merge policy. The encyclopedia has arbitration and protected pages. Every enduring commons in Ostrom's sample had rules, monitors and sanctions. A "medium without governance" is an unmanaged open-access resource, which is the case the tragedy argument correctly describes.

It does not remove cost; it changes its shape. Coordinator cost is concentrated, visible, and budgeted. Medium cost is distributed, invisible, and unbudgeted — it appears as everyone spending some minutes maintaining shared state. Distributed invisible cost is exactly the kind that goes unmeasured and therefore unfunded, which is Star & Ruhleder's problem (Chapter 23) arriving early.

It is slower for genuinely novel decisions. Stigmergy excels at incremental work against an existing structure. It is poor at the standing start — a decision with no precedent, no existing state to modify, and a deadline. For that class of decision a temporary centre is often correct, and the design question is whether it dissolves afterwards.

It requires a medium that can represent the work. Some work does not deposit a legible trace: a difficult conversation, a judgement call, an intuition about a person. Where the trace cannot be made, the mechanism does not apply, and pretending otherwise produces a medium full of ritual entries that nobody reads. **If the trace cannot be honestly deposited, do not build the medium for that class of work.**

The Specification

S21.1 — Medium requirements (VWDLA)

A shared medium **shall** satisfy all five, each scored 0–4 by the inspection tests of §21.4. Score >= 3 on all five to pass; any single score of 0 or 1 is a blocking defect.

#	Property	Requirement	Test	Pass
1	**V**isible	Current state readable by any participant without asking a person	Three-participant sixty-second test	median <= 60 s, 0 asks
2	**W**ritable	<= 2 non-blocking steps from observation to recorded state	Step count	<= 2
3	**D**urable	No critical state at bus-count 1; entries intelligible after 14 days	Bus-count census; fortnight test	0 items at bus-count 1
4	**L**ow friction	One typical write completes end to end in under 30 s	Stopwatch, 5 samples, median	median < 30 s
5	**A**ppend-oriented	State as of any past date reconstructible	Reconstruct a named date last quarter	success

S21.2 — The medium shall display its own unfinished edges

There **shall** exist a standing, visible representation of what is incomplete: open items, failing checks, known gaps, unanswered questions. **A medium that shows only finished states sup-**

plies no stimulus and is not stigmergic. Requirement: the incomplete-edge view is reachable in one step from the medium's default view.

S21.3 — Ostrom conformance checklist

Score each 0 (absent) / 1 (partial) / 2 (present). This is the governance audit and it is mandatory before any centreless design is declared complete.

#	Principle	Evidence required	Score
P1	Clearly defined boundaries	A written list of who may withdraw, and a written definition of the resource	/2
P2	Congruence: rules to conditions, appropriation to provision	Rules reference actual local conditions; take is tied to contribution	/2
P3	Collective-choice arrangements	A named, used procedure by which affected users change operational rules	/2
P4	Monitoring	Monitors are the appropriators or accountable to them; monitoring is routine, not exceptional	/2
P5	Graduated sanctions	At least three distinct severity levels, with the mildest actually used in the last year	/2
P6	Conflict-resolution mechanisms	A local arena, resolution time measured in days, no external cost	/2
P7	Right to organise recognised	No external authority has voided a local rule in the review period	/2
P8	Nested enterprises (if part of a larger system)	Each layer's remit written; cell size within the headroom-positive band	/2

Total /16. Below 10, the arrangement is not a governed commons and the tragedy argument applies to it as written.

S21.4 — Cell sizing

Working cells **shall** be sized within the headroom-positive band computed for the system's own `eta` and per-path cost. For the reference parameters of §21.5 (`eta = 0.5`, per-path cost 0.02) the band is `n <= 11` with peak at `n = 5`. Each installation **shall** recompute the band with its own measured per-path cost rather than adopting these figures.

S21.5 — Focal-point inventory

Before building any medium, the designer **shall** produce a written inventory of at least five focal points already available to participants (§21.3), and **shall** adopt them as defaults. Every adopted focal point is a `C` reduction obtained at zero build cost.

S21.6 — Hub census

The designer **shall** record, for every class of critical state, the number of people who could reconstruct it from what exists outside their own memory. **Any class with a count of 1 is a blocking defect** and must be remedied by depositing that state in the medium before the design is signed off.

S21.7 — Deposit rule

Every participant-facing procedure **shall** specify, in one sentence, what it deposits in the medium. **A procedure that deposits nothing has `eta = 0` and shall be marked as such in the survey.** This applies to meetings, reviews, inspections, and conversations held for the purpose of alignment.

S21.8 — Ethics gate (CANON §5.4)

No mechanism in this chapter may be implemented in a form that removes a participant's ability to decline. Specifically: monitoring under P4 **shall** be mutual and peer-accountable, not unilateral; participation under P3 **shall** be an available right and not a compulsory duty; and withdrawal from the commons **shall** remain possible on stated terms. **A monitoring regime whose subjects cannot opt out of the commons is not Ostrom's Principle 4; it is surveillance, and it is out of scope for this manual.**

The Failure Modes

FM-21.1 — The Phantom Centre

Description. The organisation has declared itself decentralised, removed the coordinator's title, and kept the coordination. State still lives in one person's head; they simply no longer have the authority or the budget that used to come with holding it.

Damages. `C` (unbudgeted), `delta` (unchanged and high), `eta` (falls — the person now has no mandate to convert).

Diagnostic signature. The hub census returns a bus-count of 1 for most critical state despite an org chart with no centre. People still ask the same one person. Their calendar is the busiest in the group.

Fix. Either re-establish the centre honestly and budget it, or execute S21.6 properly: enumerate what that person holds and deposit each item in the medium, one at a time, until the bus-count census passes. There is no third option.

FM-21.2 — Broadcast Mistaken for Medium

Description. A read-only dashboard, status report or weekly digest is presented as the shared state. Participants can see it; they cannot write to it.

Damages. `eta` (near zero for all receivers), `u` (falls — observations that cannot be deposited are eventually not made).

Diagnostic signature. Fails S21.1 test 2 outright. The tell is that corrections are sent to a person by message rather than made to the artefact.

Fix. Make the artefact writable with after-the-fact review. If it genuinely cannot be made writable, stop calling it the shared state and identify what the actual shared state is — usually a chat thread — and apply VWDLA to that instead.

FM-21.3 — Ritual Deposit

Description. Everyone writes to the medium because it is required. Nobody reads it. The entries are formally complete and informationally empty.

Damages. `C` (pure dissipation — everyone pays the write cost, nobody collects the read benefit), `eta` (zero despite apparent activity).

Diagnostic signature. High write volume, near-zero read volume. Ask three participants what changed in the medium last week; if none can answer, the deposits are ritual. Entries are uniform in length and phrasing, which is the fingerprint of a required field being satisfied rather than a state being recorded.

Fix. Delete required fields until writes become informative. Then make one visible consumer of the medium — a decision, a schedule, a rota — depend on it, so that a wrong deposit produces a visible wrong outcome. **A medium nothing depends on will be filled in wrongly forever and nobody will ever find out.**

FM-21.4 — Free-Ride Drift

Description. Contribution slowly concentrates in a shrinking fraction of participants. Nothing breaks. The contributors become quietly bitter and then leave, at which point something breaks all at once.

Damages. `C` (rises with `phi`), then `Q` (step loss when contributors depart).

Diagnostic signature. Plot contribution share by participant over time. Watch for the top-quartile share rising quarter on quarter. **This is measurable before it is felt**, which is the entire reason to measure it.

Fix. Ostrom P4 and P5 together: make contribution visible (not ranked, not gamified — merely visible), and install a graduated response that begins with a private word. The arithmetic of §21.6 shows monitoring pays for itself many times over. Do not begin with the severe sanction; a severe sanction alone is FM-21.5.

FM-21.5 — Terminal Sanction Only

Description. The only available response to rule-breaking is expulsion or an equivalent. It is therefore never used. The rules become advisory and everybody knows it.

Damages. `C` (rises — enforcement is impossible, so violations accumulate), `Q` (rules that are not enforced are not stored order; they are stored text).

Diagnostic signature. Ask when the sanction was last applied. If the answer is "never" or "once, years ago, and it was very ugly," you have this fault. Count the severity levels: one level is the fault.

Fix. Build at least three graduated levels and use the mildest deliberately within the first month, so that its use becomes ordinary rather than an event.

FM-21.6 — Medium Capture

Description. Access controls, mandatory fields, approval workflows and schema requirements accumulate on the medium — each individually justified — until reading requires permission and writing requires an intermediary. The medium has become a coordinator with a database interface.

Damages. `C` (rises), `eta` (falls), `u` (falls — the schema now dictates what can be said).

Diagnostic signature. Compare today's V and L scores against the scores at installation. **Medium capture is only visible as a trend, which is why S21.1 must be re-run annually and the scores retained.** A second tell: a shadow medium has appeared — a chat channel or a spreadsheet where the real state now lives.

Fix. Schema and permission budget. Fix the maximum number of required fields and the maximum number of access tiers at installation; adding one requires removing one. The shadow medium is not the problem, it is the diagnosis: whatever people fled to is telling you which property you lost.

FM-21.7 — Centreless and Ungoverned

Description. The medium exists, VWDLA passes, and the Ostrom score is 4/16. Boundaries are vague, nobody monitors, there are no sanctions, and disputes go to whoever shouts.

Damages. All of them, eventually, via free riding and boundary erosion.

Diagnostic signature. S21.3 total below 10 with a good S21.1 score. This combination is common in new, enthusiastic, technically competent groups, and it is the standard prelude to the collapse that later gets described as "we needed more structure all along."

Fix. Install P1, P4, P5 and P6 first — boundaries, monitoring, graduated sanctions, and a cheap local arena. These four cover the bulk of the failure surface. P2, P3 and P8 can follow.

FM-21.8 — Nesting Refused

Description. The body grows past the headroom-positive band and refuses to split, usually for reasons of unity, culture or fear of fragmentation.

Damages. `C` (quadratic), `h` (negative), `u` (falls — large bodies suppress dissent by fatigue long before they suppress it by policy).

Diagnostic signature. Compute `n(n−1)/2 * (per−path cost)` against `eta ln N`. If `C` exceeds `eta ln N − delta Q`, Corollary 3.4 applies and **no arrangement of the people prevents decline.** Observable proxies: meeting length rising with no rise in decisions; a growing tier of people who attend but do not speak; decisions being made before the meeting and ratified in it.

Fix. Ostrom P8. Split into cells within the band, write each cell's remit, and give the federation layer only what cells genuinely cannot do. Expect the split to be resisted on grounds of culture; the arithmetic does not care about culture.

FM-21.9 — Evaporation Refused

Description. Nothing in the medium ever expires. The trail to the exhausted food source is maintained in perpetuity. New participants cannot tell live state from archaeology.

Damages. `C` (search cost rises), `eta` (falls — the medium's signal-to-noise degrades), and eventually `h` via the maintenance term (Chapter 23).

Diagnostic signature. Ask a participant of under six months' standing to identify the current state of three items. Count their errors caused by stale entries. **More than one error in three is the fault.**

Fix. Deliberate `delta` — the pheromone must evaporate. Expiry dates on entries, an archive tier that is retrievable but out of the default view, and a standing review that retires rather than accumulates. This is Chapter 23's shedding operation applied at the level of the medium.

Worked Example — The Ashfield Tool Library

A neighbourhood tool library: 48 member households, a shared inventory of hand and power tools, a converted garage, and a shared obligation to maintain both. It is a textbook common-pool resource — subtractable (one person has the mitre saw, another does not) and hard to exclude from.

1. Encoding

Following Paper E §5.1, the five operational choices are fixed and stated before any number is computed.

Choice	Setting
State encoder	Each member household's activity assigned to one of N = 8 craft categories: woodwork, metalwork, garden, automotive, textiles, electronics, masonry, cycle repair
Interaction observable	Cross-category tool loans and joint working sessions per month
Order observable Y	Composite: tools in working order and correctly located, live maintenance skills held by more than one member, documented repair procedures
Time scale	One month = one period
Cost observable	Volunteer hours plus consumables plus loss and breakage, normalised to the same units as Y

The Y justification required by Paper E §6.3 — that the order proxy is not a repackaging of the input — holds here: Y is measured from a physical inventory audit and a skills census conducted independently of the loan log that produces U. A month with heavy borrowing and no repairs scores high U and flat Y. **The two observables can and do move in opposite directions, which is the test.**

2. Measured parameters

N = 8, so ln N = 2.0794.

Non-parallelism from the disagreement-and-decorrelation protocol of Chapter 9: u = 0.62.

```
N_eff = 1 + (N−1)u = 1 + 7(0.62) = 5.34
S     = ln(N_eff) = ln(5.34) = 1.6752
```

Forty-eight households, eight categories, and an effective count of 5.34. The audit for redundancy removed a quarter of the categorical variety before any surplus was credited.

3. Regime A — the coordinator

The library's first two years ran on a part-time paid coordinator who held the loan ledger, the maintenance schedule, the key rota and the membership list.

```
eta   = 0.35     interaction converts only via the coordinator's records
delta = 0.25     the coordinator's own turnover and forgetting rate
C     = 0.52     stipend, scheduling, queueing on one person's availability
```

Headroom at the start, Q_0 = 0:

```
h_0 = eta ln N − delta Q_0 − C
    = 0.35(2.0794) − 0 − 0.52
    = 0.7278 − 0.52
    = +0.2078
```

Positive but thin. Steady state by Theorem 3.5:

```
eta S = 0.35 x 1.6752 = 0.5863
Q_inf = (eta S − C) / delta = (0.5863 − 0.52) / 0.25 = 0.2653
```

Trajectory from `Q_0 = 0`:

t (months)	0	2	4	6	8	10
Q_t	0.000	0.116	0.181	0.218	0.239	0.250

The library reaches a steady state of `Q_inf = 0.265` and stays there. It works. It never gets better. Check the derived identity of §21.0.2:

```
h at Q_inf = 0.35(2.0794) − 0.25(0.2653) − 0.52 = +0.1415
eta ln(N/N_eff) = 0.35 x (2.0794 − 1.6752) = 0.35 x 0.4042 = +0.1415    [agrees]
```

Then the coordinator moves house. Their successor inherits the ledger but not the relationships, not the knowledge of which drill has the intermittent fault, not the informal understandings about who gets priority in spring. The hub census had a bus-count of 1 on most critical state. Estimated loss: 70 per cent of `Q` in a single step.

```
Q drops from 0.239 to 0.072
```

Recovery to the previous level takes approximately six months. **The library did not lose a volunteer. It lost most of its stored order, because that is where the order was stored.**

4. Regime B — the medium

The redesign. State is moved out of the coordinator and into two physical artefacts and one digital one:

- A **peg board** with a hook per tool: the tool's presence or absence, the borrower's tag, and the return date, all readable from the doorway in three seconds. Visible: 4/4.
- A **bound logbook** on a chain by the door: date, tool, member, condition on return, one line for anything odd. Writable in under fifteen seconds with a pen. Writable 4/4, low-friction 4/4, append-oriented 4/4.
- A **shared spreadsheet** mirrored weekly from the logbook by whoever is on rota, holding the membership list, maintenance schedule and repair procedures. Durable 3/4.

VWDLA total: 19/20. Passes S21.1.

The unfinished-edge display required by S21.2 is a red tag hung on the peg for any tool needing attention, plus a chalked "needs doing" list on the garage wall. **The board displays its own defects, which is what makes it a stimulus rather than a report.**

Revised parameters, measured over the following year:

```
eta   = 0.55    every loan and every fault deposits directly into durable record
delta = 0.12    state survives departures; loss is physical wear and slow drift
C     = 0.45    no stipend; rota cost and the write-cost distributed across members
```

Headroom at `Q_0 = 0`:

```
h_0 = 0.55(2.0794) − 0 − 0.45 = 1.1437 − 0.45 = +0.6937
```

Steady state:

```
eta S = 0.55 x 1.6752 = 0.9214
Q_inf = (0.9214 − 0.45) / 0.12 = 3.9281
```

Trajectory:

t (months)	0	2	4	6	8	10	12
Q_t	0.000	0.886	1.572	2.104	2.515	2.834	3.081
h_t	+0.694	+0.587	+0.505	+0.441	+0.392	+0.354	+0.324

Identity check again:

```
h at Q_inf = 0.55(2.0794) − 0.12(3.9281) − 0.45 = +0.2223
eta ln(N/N_eff) = 0.55 x 0.4042 = +0.2223     [agrees]
```

Time constant `tau = 1/delta = 8.33` months; half-life of stored order under zero input is `ln 2 / (−ln(1−0.12)) = 5.42` months.

Regime B holds 14.8 times the stored order of Regime A — 3.928 against 0.265 — at lower total cost. The mechanism is not effort. Nobody in Regime B works harder than the coordinator did. The mechanism is that `eta` rose from 0.35 to 0.55 because every act of borrowing now deposits, and `delta` fell from 0.25 to 0.12 because the deposit is not in a person.

5. The turnover test, repeated

Regime B's most active member also leaves, at month 8. The hub census now shows no critical state at bus-count 1; what is lost is relationship and judgement, not record. Estimated loss: 15 per cent.

```
Q drops from 2.515 to 2.138
```

Recovery: under one month. Compare Regime A's 70 per cent loss and six-month recovery. **Same event, same organisation, same people. Different location for the state.**

6. Free riding and the monitoring decision

Twenty months in, the contribution-share plot required by FM-21.4 shows the top quartile of households supplying 71 per cent of maintenance hours, up from 58 per cent. Estimated free-riding fraction `phi = 0.35`, loss coefficient `lambda = 0.9`, base cost 0.40:

```
C     = 0.40 + 0.9(0.35) = 0.7150
h(Q=0)= 0.55(2.0794) − 0.7150 = +0.4287
Q_inf = (0.9214 − 0.7150) / 0.12 = 1.7198
```

Stored order has been more than halved by free riding — 3.93 down to 1.72 — without a single visible failure. Nothing broke. It simply costs more than it did.

Install Ostrom P4 and P5. **Monitoring**: the logbook gains a column for maintenance hours, and the monthly rota is posted on the same board as the tools — mutual, peer-accountable, a by-product of work already being done, satisfying S21.8. **Graduated sanctions**: three levels, agreed by the members under P3 — (1) a private word from the rota partner; (2) loss of priority booking for one month; (3) suspension of borrowing rights pending a meeting under P6.

Monitoring cost 0.08. Measured `phi` after two quarters: 0.06.

```
C     = 0.40 + 0.08 + 0.9(0.06) = 0.5340
h(Q=0)= 0.55(2.0794) − 0.5340 = +0.6097
Q_inf = (0.9214 − 0.5340) / 0.12 = 3.2281
```

Eight hundredths of a unit of monitoring cost bought back 1.51 units of steady-state stored order — an 88 per cent recovery of what free riding had taken. Level 1 of the sanction ladder was used eleven times in the first year; level 2 twice; level 3 never. That distribution is the signature of a working graduated system, and it is the direct answer to FM-21.5.

7. Ostrom conformance audit

Principle	Evidence at Ashfield	Score
P1 boundaries	Written membership roll; inventory list defines the resource	2
P2 congruence	Borrowing limits set by tool scarcity; borrowing rights tied to rota participation	2
P3 collective choice	Rules changed at the quarterly members' meeting; sanction ladder was written there	2
P4 monitoring	Peg board and logbook are mutual and by-product; no external monitor	2
P5 graduated sanctions	Three levels; mildest used 11 times this year	2
P6 conflict resolution	Any two members may call a meeting within seven days; no cost	2
P7 right to organise	Independent association; no external authority has voided a rule	2
P8 nesting	Four craft circles of 8–14 households, plus a delegate council of 4	1

Total: 15/16. P8 scores 1 because two of the four craft circles exceed the headroom-positive band computed with Ashfield's own per-path cost, and the delegate council's remit is not written down. Both are on the remediation list.

8. What would falsify this

Stated in advance, per Chapter 25:

1. If, after installing the medium, measured `delta` did not fall below 0.20, **Postulate M1 failed in this system** and the redesign should be reported as unsuccessful regardless of how good the board looks.
2. If `Q` under Regime B did not exceed Regime A's steady state within 12 periods, the modelling of `eta` is wrong.
3. If the monitoring intervention did not reduce measured `phi` within two quarters, Ostrom P4 does not transfer to this setting, and that should be reported as a negative result, not explained away.
4. The shuffled-input control of Chapter 25 must be run: if `Y` predicts as well from shuffled `U` as from real `U`, the whole model is fitting a trend and nothing more.

9. The rule, extracted

Ashfield's redesign changed no personnel, no funding, and no rules of membership. It changed where the state lived. That single move raised `eta` by 57 per cent, cut `delta` by 52 per cent, and multiplied steady-state stored order fifteenfold. Then a monitoring regime costing 0.08 units recovered most of what free riding had quietly removed.

Put the state in the medium, not in a coordinator. Then govern the medium by Ostrom's eight, or the tragedy argument applies to you exactly as written.

Chapter 22 . Individuality as Load-Bearing Structure

The Law

This is the ethical chapter of the manual. It is also the mathematical one, and that is its entire force.

Every other argument for preserving individuality in a group is a moral argument, and moral arguments lose to schedule pressure. The manager who wants the meeting over, the committee that wants a clean recommendation, the department that wants alignment — none of them are villains, and none of them will be moved by an appeal to the dignity of the dissenter. They will be moved, if at all, by the demonstration that **the thing they are trying to produce cannot be produced by the method they are using.**

That demonstration exists. It arrives from four independent directions, in four separate literatures, using four unrelated formalisms, and it says one thing.

22.0.1 What is proven

From Paper A, Theorem 4.3 (extremality), stated exactly:

```
F = 0        if and only if the states are parallel
F = ln N     if and only if the states are orthogonal
```

And from the effective-count identity:

```
N_eff = 1 + (N − 1) u          u = 1 − (x . y)^2 = || x ^ y ||^2   in [0,1]
S     = ln N_eff
```

The consequence is arithmetical and not negotiable. **When the members of a group become parallel, `u` goes to 0, `N_eff` goes to 1, and the surplus of their interaction is `ln 1 = 0` — exactly zero, regardless of how many of them there are, how able they are, or how hard they work.**

Nine people who have converged are not nine people. They are one person, counted nine times, at nine times the cost.

This is a theorem about a mathematical object, and Paper A makes no empirical or metaphysical claim (CANON §1.1). The empirical content of this chapter comes from elsewhere — from Condorcet, Hong & Page, Bikhchandani/Hirshleifer/Welch, and Woolley — and the surplus law's contribution is that it prices the same phenomenon in the same units as everything else in this manual.

22.0.2 The four results

Source	Formalism	Statement	What it says about individuality
Condorcet (1785)	Probability	Majority accuracy exceeds individual accuracy and tends to 1 as n grows — **only under independence**	Independence is a *precondition*, not a courtesy
Hong & Page, *PNAS* 2004	Computational / analytic	Under stated conditions, a diverse group of problem solvers outperforms a group of the individually best	Difference contributes something ability cannot
Bikhchandani, Hirshleifer & Welch (1992)	Sequential Bayesian	Rational observation of others' choices produces cascades in which private information stops entering the pool	Copying is individually rational and collectively destructive
Woolley et al., *Science* 2010	Factor analysis	A collective-intelligence factor c correlates with equality of conversational turn-taking, not with member IQ	The *distribution* of participation is the driver, not the peaks

Four methods. One conclusion. **Preserving individuality is not a concession to comfort. It is a requirement for the collective to function at all.**

22.0.3 The postulate, badged

Postulate I1 (the merge). Human groups under conditions of shared training, shared incentive, prolonged co-location and social pressure drift toward parallelism — u declines monotonically absent maintenance. This is a **design postulate**, not a theorem and not a measured universal. It is consistent with the cascade literature and with the general phenomenon Chapter 11 catalogues, but this manual has not measured a drift rate, and any installation that needs one must measure its own by the protocol in §22.7.

What *is* derivable, given I1, is the consequence: u is not a property a group acquires. **It is a quantity that decays, and therefore it is a `delta` problem as much as a design problem.** A group that was diverse in January and was not maintained is not diverse in December, and nobody will have noticed, because the experience of convergence from the inside is indistinguishable from the experience of being right.

22.0.4 The ethical boundary, stated before the engineering

CANON §5.4 is binding on every protocol in this chapter: **no intervention that removes a person's ability to decline is in scope.** This is not a caveat appended for politeness; it is a design constraint that determines which protocols appear below and which do not.

The constraint has teeth. It rules out compulsory disclosure of private estimates. It rules out assigning a person the permanent role of dissenter without their consent. It rules out any "diversity" measure that conscripts a person's identity as an instrument. It rules out mandatory participation in elicitation. Every protocol in §22.7 is specified with an explicit decline path, and where a protocol cannot be given one, it is excluded from the manual and the exclusion is stated.

There is a happy structural fact here, and it is worth naming because it recurs throughout Part IV. **The protocols that preserve individuality work by removing pressure, not by applying it.** Their worst failure mode is that someone declines to participate and the group learns slightly less than it might have. They have no coercive failure mode, because coercion is not their mechanism. That is the same safety property Chapter 20 identifies for the silo intervention, and it is not an accident: methods that operate by lowering resistance cannot be weaponised in the way methods that operate by raising it can.

The Engineering

22.1 The arithmetic of the merge

Start with the price list, because everything downstream is an argument about how to avoid paying it. A nine-person panel, `N = 9`, `ln N = 2.1972`:

u	N_eff	S = ln N_eff	Fraction of ceiling
0.00	1.00	0.0000	0%
0.10	1.80	0.5878	27%
0.12	1.96	0.6729	31%
0.25	3.00	1.0986	50%
0.50	5.00	1.6094	73%
0.69	6.52	1.8749	85%
0.75	7.00	1.9459	89%
1.00	9.00	2.1972	100%

Two features of this table are worth dwelling on.

First, the collapse at the left end is total, not gradual in effect. At `u = 0` the surplus is not small; it is zero. The panel is a single instrument, and its nine readings are one reading. The law is not being poetic here: `ln(1) = 0` exactly.

Second, the recovery is front-loaded. Going from `u = 0` to `u = 0.25` — a quarter of the way to orthogonality — buys half the available surplus. This is Paper A Theorem 3.3, strict concavity, arriving as a design fact: the first increment of genuine difference is worth more than all subsequent increments combined, which means the first honest dissenter in a converged room is the single highest-value intervention available in that room.

The same concavity governs headcount. At `u = 0.62`, the marginal surplus of the N-th genuinely distinct member:

N	2	3	4	5	6	8	10
S	0.4824	0.8065	1.0508	1.2470	1.4110	1.6752	1.8840
marginal	+0.4824	+0.3240	+0.2443	+0.1962	+0.1640	+0.1234	+0.0990

The second genuinely different person is worth roughly five times the tenth. Concavity is merciful — it means you do not need a cast of thousands — and it is unforgiving in the other direction, because it also means that adding people who are *not* different adds nothing at all while adding cost quadratically (Chapter 5).

22.2 Condorcet: independence is the load-bearing clause

The Marquis de Condorcet published his jury theorem in 1785, in *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. The setup:

- A binary question with a correct answer.
- `n` voters, `n` odd.
- Each voter is correct with probability `p`, **independently of the others**.
- The group decides by simple majority.

The theorem: if `p > 1/2`, the probability that the majority is correct exceeds `p`, increases monotonically in `n`, and tends to 1 as `n` tends to infinity.

Compute it. The probability the majority of `n` is correct is

```
P(n, p) = sum over k from (n+1)/2 to n of  C(n,k) p^k (1−p)^(n−k)
```

n	p = 0.55	p = 0.60	p = 0.70
1	0.5500	0.6000	0.7000
3	0.5748	0.6480	0.7840
5	0.5931	0.6826	0.8369
9	0.6214	0.7334	0.9012
15	0.6535	0.7869	0.9500
25	0.6937	0.8462	0.9825
51	0.7641	0.9265	0.9986
99	0.8413	0.9781	1.0000

Ninety-nine people each barely better than a coin — 55 per cent — reach 84 per cent as a body. This is the most encouraging table in the sociology of groups and it is why democratic aggregation is not merely a value but a technique.

Now the clause everybody quotes past. Read the theorem again: *independently of the others*. Every number in that table is conditional on it. And the theorem has a second half that is almost never taught:

If `p < 1/2`, the same monotonicity runs the other way. The majority is *worse* than the individual, and gets worse with `n`:

n	1	3	9	25	99
P at p = 0.45	0.4500	0.4253	0.3786	0.3063	0.1587

Ninety-nine people who are each 45 per cent accurate produce a body that is **16 per cent accurate.** Aggregation is an amplifier. It amplifies whatever sign the individual competence has. And at exactly `p = 0.5`, `P(n, 0.5) = 0.5` for every `n` — no amount of aggregation extracts signal from noise.

What correlation does Independence is an idealisation. Real voters share training, share news sources, share incentives, and talk to each other. Model this with the simplest possible correlation structure, **badged as a modelling choice, not as a result**:

Model C1 (shared-signal mixture). With probability `rho`, all voters observe and follow a single common signal, itself correct with probability `p`. With probability `1 − rho`, they vote independently at accuracy `p`. Then

```
P(n, p, rho) = rho * p + (1 - rho) * P_independent(n, p)
```

This is crude — it is a two-point mixture, not a general correlation model — but it is transparent and it is enough to make the point. With `n = 9`, `p = 0.60`, so `P_independent = 0.7334`:

rho	0.0	0.2	0.4	0.6	0.8	1.0
P	0.7334	0.7067	0.6801	0.6534	0.6267	0.6000

At `rho = 1`, the nine-person panel performs exactly as well as one person, because it *is* one person. That is Theorem 4.3 again, in probability rather than geometry: **`rho -> 1` is `u -> 0` is `N_eff -> 1` is `S -> 0`.**

And scale does not rescue it. Ninety-nine voters at `p = 0.60` with `rho = 0.30` give 0.8646, against 0.9781 under independence. **Correlation is not a discount on the benefit of scale; it is a ceiling on it.** Adding voters to a correlated pool buys progressively less and eventually nothing, exactly as adding parallel windows to a wall buys no view.

Honesty note on the bridge. `rho` and `u` are measures of the same underlying thing — non-redundancy — in different formalisms, and this manual does not claim a general identity between them. Converting one into the other requires a specific state encoder (Paper E §5.1) and the conversion should be stated whenever it is made. What is claimed is the direction: **as correlation rises, effective count falls, by both accountings.**

22.3 Hong & Page: diversity trumps ability

Lu Hong and Scott Page published "Groups of diverse problem solvers can outperform groups of high-ability problem solvers" in *PNAS* in 2004. The framework: problem solvers search a solution landscape using *heuristics* — internal representations of the problem plus rules for where to look next. A solver gets stuck at a point that is a local optimum *for that solver's heuristic*. A group works by passing the current best solution around until no member can improve it.

The result, informally: under stated conditions, a group of randomly selected problem solvers drawn from a large pool of able solvers outperforms a group formed by taking the individually highest-performing solvers.

The mechanism is entirely legible in this manual's terms. **A solver's ability determines how good their local optimum is. A solver's difference determines whether they are stuck at the same place as everyone else.** A group of the individually best solvers tends to share heuristics — that is frequently *why* they are the best on the measured task — and a group that shares heuristics shares local optima. When the group gets stuck, everyone is stuck at the same point, and there is nobody whose different representation makes the next move visible.

The conditions matter and this manual states them rather than eliding them. The result holds when: the problem is difficult (no solver alone finds the global optimum); the solvers are individually capable (their local optima are not trivially bad); the pool from which the group is drawn is large; and the group is large enough that random selection produces genuine heuristic variety.

Honesty note. The Hong–Page theorem has been the subject of substantive mathematical criticism — most prominently Abbott Thompson's 2014 critique in the *Notices of the American Mathematical Society*, which argues that the formal result depends heavily on the specific construction of "diversity" and does not support the broad interpretation often placed on it. This manual's position is narrow and, it believes, uncontested by that critique: **the mechanism — that shared heuristics produce shared local optima, and that a group stuck in one place needs someone standing somewhere else — does not require the strong version of the theorem.** It requires only that different representations get stuck in different places, which is nearly a definition. Where this manual uses Hong & Page, it uses the mechanism, and the reader is directed to the primary literature and the critique for the formal claim.

22.4 Cascades: how copying destroys the information it appears to use

This is the most important result in the chapter, because it explains why u collapses in groups that nobody is pressuring, composed of people who are individually behaving well.

Bikhchandani, Hirshleifer & Welch (1992) analysed sequential decision-making under observation. Agents decide one at a time. Each has a private signal about which of two options is correct. Each can observe the *actions* of everyone before them, but not their signals.

Their finding: after a very small number of decisions, agents rationally stop using their private signals. An **informational cascade** forms — everyone thereafter copies the visible pattern regardless of what they privately know — and **from that moment on, no further private information enters the public pool.**

The chain looks informative. It is not. It contains, permanently, the information of the first two or three agents and nothing more.

The arithmetic Take the standard symmetric setup. The true state is HIGH (so ADOPT is correct). Each agent receives a private signal, correct with probability `p > 1/2`. Agents act in sequence; ties are broken by a fair coin.

- **Agent 1** has only their own signal, so adopts with probability `p`.

- **Agent 2** observes agent 1's action and their own signal. If the signals agree, they follow. If they disagree, the posterior is exactly balanced and they flip a coin. So given agent 1 adopted, agent 2 adopts with probability `p + (1−p)/2 = (1+p)/2`.
- **Agent 3** observes two actions. If both were the same, agent 3's single private signal cannot overturn two, so agent 3 copies regardless of their own signal — **a cascade.** If the first two actions differed, they cancel, and agent 3 is in agent 1's position: the process restarts.

The probabilities:

```
P(up cascade starting)   = p (1+p) / 2
P(down cascade starting) = (1−p)(2−p) / 2
P(restart)               = 1 − [p(1+p)/2 + (1−p)(2−p)/2] = p(1−p)
```

Since a cascade eventually forms with probability 1, the probability that the cascade that forms is the *correct* one is

```
P(correct cascade) = [ p(1+p)/2 ] / [ p^2 − p + 1 ]
```

Evaluated:

p	0.55	0.60	0.62	0.70	0.80	0.90
P(correct cascade)	0.5664	0.6316	0.6570	0.7532	0.8571	0.9396

Verification. A Monte Carlo simulation of this process at `p = 0.60` over 400,000 trials returned 0.6323 against the analytic 0.6316.

Read the comparison and do not look away from it At `p = 0.60`:

Arrangement	Accuracy
One person alone	0.6000
An unbounded cascade — a million observers	**0.6316**
9 independent voters, majority rule	0.7334
99 independent voters, majority rule	0.9781

An infinite queue of rational, well-informed, honest people, each holding genuine private information, performs worse than nine of them who never looked at each other. The million observers contributed nothing. Their information exists; it simply never entered the pool, because each of them rationally deferred to a public record that contained only the first two people's signals.

There is no villain in this story. No one is lazy, conformist, or cowardly. Every agent behaves optimally given what they can see. **The failure is structural: the observation of others' actions before the expression of one's own signal is sufficient to destroy the aggregation, all by itself.**

That is the design brief. If observation-before-expression destroys aggregation, then the fix is to **move expression before observation.** Everything in §22.7 is a variation on that single move.

22.5 Woolley: the distribution, not the peak

Woolley, Chabris, Pentland, Hashmi & Malone (*Science*, 2010) had groups perform a battery of varied tasks and found that performance across the battery loaded onto a single statistical factor, `c`, analogous to the individual `g`. The findings that matter here, stated precisely:

- `c` was **not significantly predicted by the average member intelligence, nor by the maximum**.
- `c` correlated with average **social sensitivity** of members, measured by the Reading the Mind in the Eyes test.
- `c` correlated with **equality in the distribution of conversational turn-taking**. Groups in which a few members dominated the conversation had lower `c` than groups in which speaking was distributed.
- `c` correlated with the **proportion of women** in the group, an effect largely mediated by social sensitivity.

The turn-taking finding is the one this chapter builds on. Consider what unequal turn-taking *is*, in the terms of §22.1. A group in which two people speak 80 per cent of the time is a group in which seven people's states are not entering the interaction at all. The nominal `N` is nine; the participating `N` is two. **Turn-taking inequality is `N_eff` suppression by airtime**, and it does not require anyone to be silenced — the ordinary dynamics of a fast, unstructured conversation are sufficient.

Honest limits, per CANON. Woolley's `c` was measured in short-task laboratory settings with ad hoc groups. Extrapolation to standing institutional bodies is a **modelling assumption** and is badged as such throughout this manual. What transfers with confidence is the narrower claim: within the studied setting, the distribution of participation predicted performance and the peak of member ability did not.

22.6 Four results, one statement

Lay them side by side.

```
Condorcet    : aggregation works only if the inputs are independent
Hong & Page  : difference contributes what ability cannot
BHW 1992     : copying stops information from entering the pool
Woolley 2010 : equal participation predicts collective performance
```

Each is a statement about a different thing — probability, search, sequential inference, factor structure. Each was produced by people not working on the others' problems. And each reduces to the same instruction:

> **Keep the individuals individual, and keep them all in.**

The surplus law prices it. `S = ln(1 + (N−1)u)`. Merge the members and `u -> 0`, `N_eff -> 1`, `S -> 0`. Silence the members and the participating `N` falls, with the same effect. **Merge the individuals and the c-factor falls, the cascade sets in, and the group becomes one mind with one blind spot** — a blind spot that is now nine times as confident as it was, because nine people agree.

The confidence is the cruellest part. A converged group does not experience itself as impoverished. It experiences itself as *aligned*, and alignment feels exactly like being right. There is no

internal signal that distinguishes a room that has converged on the truth from a room that has converged on the first speaker. **Only a protocol can distinguish them, and only if the protocol ran before the room started talking.**

22.7 The protocols

Five. Each is specified with its mechanism, its cost, its expected effect, and — mandatorily — its decline path.

P-1 . Independent estimate before discussion Mechanism. Before any discussion of the item, every participant privately records their own judgement — an estimate, a rating, a recommendation, a probability — with a one-line reason. The records are sealed until all are in. Only then does discussion open, and the distribution of independent judgements is displayed *first*.

Why it works. It is the direct structural answer to §22.4. The cascade requires observation before expression; this protocol inverts the order. Every private signal is entered into the pool before any public signal exists to override it. **It is not a debiasing technique and does not depend on anyone being unbiased. It changes the order of operations.**

Cost. Four to eight minutes per item. In the storage law this is a `C` increment; the Worked Example computes it as 0.03.

Effect. The distribution itself is the deliverable, and it is worth as much as the decision. A tight distribution before discussion is either genuine convergence or a symptom that convergence already happened elsewhere — and the two are distinguishable by asking whether the reasons given differ. **Identical conclusions with identical reasons is a `u = 0` reading. Identical conclusions with different reasons is the strongest evidence a committee can generate.**

Decline path. A participant may record "no independent view" and the item proceeds. The record of a declined estimate is itself information: a member who declines on every item in their nominal area of expertise is telling you something about the panel that no survey will.

P-2 . Anonymous-first elicitation Mechanism. The first round of judgements is collected without attribution. Attribution is added in the second round, after the anonymous distribution is visible, and only for participants who choose to attach their names.

Why it works. It removes rank from the first signal. In a group with any status gradient, the junior member's private signal is systematically discounted by its own holder before it is spoken — which is a cascade driven by hierarchy rather than by sequence, and has the same effect on the pool.

Cost. A slip of paper and a box, or a form. Effectively zero in money; three minutes in time.

Effect. The instrument to watch is the **anonymous-to-attributed drift**: the difference between the anonymous distribution and the attributed one. A large drift is a direct measurement of how much signal the group's status structure was suppressing. It is one of the few u-suppression readings that can be taken without a control group.

Decline path. Participation in the anonymous round is optional and the option is stated aloud each time, not merely in the written procedure. Anonymity that participants do not believe in is worse than no anonymity, because it produces confident false readings.

Constraint, per CANON §5.4. Anonymity must be real. If the group is small enough that handwriting, phrasing or expertise identifies the author, the protocol is not anonymous and must not be described as such. Say so and use P-1 instead.

P-3 . Dissent protection Mechanism. A written, standing, published commitment that (a) a recorded dissent carries no consequence for standing, assignment, or advancement; (b) dissents are entered in the record with the decision, not appended afterwards; (c) the person who dissents is not the person who is then assigned to implement the majority position.

Why it works. Under Postulate I1, u decays. The largest driver of decay in an institutional body is not persuasion; it is anticipation of cost. A member who expects a recorded disagreement to be remembered at review time will not record it, and their private signal never enters the pool. The protection is a `delta` intervention on u: it does not create difference, it stops the leak.

Cost. Near zero to write. Substantial to honour, and the honouring is the whole protocol. **A dissent protection that has never been tested is not known to exist.**

Effect. The measurable quantity is the **dissent rate over time**. A body whose recorded dissent rate declines quarter on quarter is a body whose u is collapsing, and the decline is visible in the minutes long before it is visible in the decisions.

Decline path. Nobody may be compelled to dissent, and — importantly — nobody may be assigned the role of standing contrarian without their explicit and revocable consent. Conscripted dissent is theatre; it produces a visible objection that carries no private signal, which is worse than silence because it looks like the protocol is working.

P-4 . Minority report requirement Mechanism. Any decision reached with less than unanimity must carry a written minority report: the strongest available statement of the rejected position, its evidence, and the specific observation that would show the majority to be wrong. The report travels with the decision permanently. Where a decision is genuinely unanimous, the requirement is met by an appointed **steelman** — a member tasked with writing the best case against, with that member's consent.

Why it works. Three effects, and the third is the largest.

1. It preserves the minority's information in the durable record. In storage-law terms, dissent that is written down has `eta > 0`; dissent that is spoken and lost has `eta = 0`. **A minority report is an `eta` device applied to the most perishable and most valuable class of interaction in the whole system.**
2. It makes the decision auditable later, when the world has supplied an answer.
3. **It converts unanimity from an achievement into an obligation.** A committee that must produce a written minority report cannot use agreement as an exit. This single change

removes most of the incentive gradient that drives premature convergence, because the cheapest path is no longer "everyone agree quickly."

Cost. Twenty to forty minutes per non-unanimous decision. Real, and the largest C item in this list.

Effect. The falsification clause in point (1) is the part usually dropped and it is the part that matters. A minority report that says "we disagree" is a mood. A minority report that says "if the programme has not produced measurable X by month 18, the majority was wrong" is an instrument, and in eighteen months it will have paid for every minute it cost.

Decline path. The steelman role is by consent and is rotated. If nobody consents, the record states that no steelman was appointed, which is itself the finding.

P-5 . Turn-taking equalisation Mechanism. Structural devices that flatten the airtime distribution without anyone policing it: round-robin on each item; the **last-speaks-first inversion**, in which the most senior person in the room speaks last; a hard first-round time box per speaker; and the physical geometry of Chapter 17, a table with no head.

Why it works. Directly targets Woolley's turn-taking finding, and simultaneously breaks rank-driven cascades: when the senior member speaks first, every subsequent contribution is a response to a public signal rather than an expression of a private one.

Cost. Near zero. Chapter 17's argument applies in full: **furniture and speaking order are the cheapest governance instruments ever built, and they enforce continuously without anyone administering them.**

Effect. Measure the **Gini coefficient of speaking time**, or more simply the share held by the top two speakers. Woolley's finding predicts that flattening this distribution raises c. Chapter 14 supplies the physical precondition — in a reverberant room the loudest voice wins by acoustics regardless of protocol, so the acoustic fix precedes the procedural one.

Decline path. A participant in a round-robin may pass. Passing is recorded as passing, not as assent, and this distinction must be explicit in the record. **A protocol that reads silence as agreement has re-created the cascade inside the mechanism designed to prevent it.**

22.8 What the protocols cost and what they buy

Costs are in the normalised units of the Worked Example, per panel item.

Protocol	C per item	Primarily raises	Primarily protects
P-1 independent estimate	0.030	u	against sequential cascade
P-2 anonymous-first	0.010	u	against rank cascade
P-3 dissent protection	0.005	u (lowers its decay)	against anticipatory self-censorship

Protocol	C per item	Primarily raises	Primarily protects
P-4 minority report	0.040	eta	against loss of the rejected case
P-5 turn-taking	0.005	participating N	against airtime concentration
Total	**0.090**		

Nine hundredths of a unit. Section 22.9 of the Worked Example shows this purchase moving u from 0.12 to 0.71 and taking the panel from a negative steady state to a positive one. **The protocols are cheap because they are procedural. They are effective because the failure they address is procedural.**

Note the split in the second column. P-1, P-2 and P-3 raise u — they get more distinct signal into the room. P-4 raises eta — it gets the signal that arrived into the durable record. **These are different parameters and neither substitutes for the other.** A panel with excellent elicitation and no minority report generates real difference every week and stores none of it. A panel with rigorous minority reports and no independent elicitation writes down, with great care, the objections of a group that has already converged.

22.9 The ethics, stated as engineering

The manual refuses to recommend any intervention that removes a person's ability to decline (CANON §5.4). It is worth being explicit about what this excludes, because the exclusions are not obvious and each has been proposed somewhere in good faith.

Excluded: compulsory disclosure of private judgement. P-1 is powerful precisely because it captures what people actually think. That power is exactly why it must not be mandatory. A required private estimate is not private, and a person who fears the record will write what is safe — which returns you to the cascade with an extra step and a false sense of rigour.

Excluded: assigned permanent dissenters. Appointing someone the standing contrarian both conscripts them and defuses them: their objection is discounted by everyone, including themselves, as role-play. The steelman under P-4 avoids this by being consensual, rotating, and about the argument rather than about the person.

Excluded: identity-as-instrument. Using a person's demographic or biographical characteristics as a proxy for their epistemic contribution is both an ethical failure and a measurement error. It is the central diagnosis of Chapter 11: demographic difference without epistemic difference gives u = 0 and buys nothing. **The quantity in the law is the non-parallelism of their states, and the only way to know it is to measure it.**

Excluded: mandatory participation. Every protocol above has a decline path, and it must be usable without explanation.

Included, and required: the right to be recorded as declining. This is the positive form of the same principle. A person who does not wish to give a view must be able to have that fact entered rather than being counted as assent. The alternative — silence read as agreement — is not a neutral default. It is a mechanism that manufactures unanimity out of absence, and a body that

uses it will report high agreement and low u, which is the exact signature of the failure this chapter exists to prevent.

The Specification

S22.1 — Order of operations

Any deliberative body **shall** collect independent individual judgements on an item **before** opening discussion of that item. **Requirement: no participant observes another participant's position on the item before recording their own.** This is the single highest-value requirement in the chapter and it is the direct structural answer to BHW 1992.

S22.2 — Distribution before conclusion

The distribution of independent judgements **shall** be displayed to the body before discussion begins, including its spread. **A body that sees only the mean has been shown a converged number and will converge on it.**

S22.3 — Reasons, not only conclusions

Each independent judgement **shall** carry at least one line of reasoning. Agreement of conclusions with agreement of reasons **shall** be recorded as a u-suppression indicator, not as corroboration.

S22.4 — Dissent record

Every non-unanimous decision **shall** carry a written minority report containing: the rejected position, its evidence, and **a specific falsifier** — the observation that would show the majority wrong. Unanimous decisions **shall** carry a consented steelman statement, or an explicit note that none was appointed.

S22.5 — Dissent protection, published and tested

A written dissent protection **shall** exist, be published, and be exercised at least once per year. **An untested protection shall be recorded in the survey as absent.**

S22.6 — Participation measurement

Airtime **shall** be measured on a sample of sessions, at least quarterly. Recorded metric: share of speaking time held by the top two speakers. Threshold for intervention: **above 60 per cent in a group of eight or more.**

S22.7 — u maintenance schedule

Because u decays (Postulate I1), the body **shall** re-measure u at least annually by the disagreement-and-error-decorrelation protocol of Chapter 9, and **shall** record the trend. A declining trend across two consecutive measurements **shall** trigger a composition review.

S22.8 — Decline paths, mandatory

Every protocol in this chapter **shall** be implemented with an explicit decline path, stated aloud at the time of use and not only in the written procedure. **Silence shall never be recorded as assent; a pass shall be recorded as a pass.** A protocol without a decline path is out of scope for this manual and shall not be installed under its authority.

S22.9 — Anonymity truthfulness

Where anonymity is claimed, the designer **shall** verify that the group size and the response format do not identify authors. Where they do, the protocol **shall** be relabelled and the anonymity claim withdrawn.

The Failure Modes

FM-22.1 — The Cascade Collapse

Description. Discussion opens with a senior or confident member stating a position. Everything subsequent is a response to that position. The panel reaches high agreement quickly and reports it as confidence.

Damages. u (to near zero within a single session), and through it S, Q.

Diagnostic signature. Order-of-speaking correlates with the final position. Time-to-consensus is short and falling. Members can state the conclusion but cannot state a reason that differs from the first speaker's.

Fix. S22.1. There is no soft version of this fix: independent recording before discussion, or the cascade continues. Debiasing training does not address it, because the agents in the BHW model are not biased.

FM-22.2 — Unanimity Reported as Strength

Description. The body's minutes show near-total agreement across all items and this is presented as evidence of quality.

Damages. u, and the credibility of the whole record.

Diagnostic signature. Dissent rate below 5 per cent of decisions over a year. **Cross-check against outcomes: a body that agrees on everything and is wrong at the base rate is a body whose agreement carries no information.**

Fix. S22.4 — the minority report requirement converts unanimity from an exit into an obligation. Expect resistance framed as efficiency; the arithmetic of §22.4 is the reply.

FM-22.3 — The Token Dissenter

Description. One member reliably objects. Everyone knows they will object. The objection is heard, noted, discounted, and the decision proceeds unchanged.

Damages. u in effect, though not on paper — the record shows dissent while the pool receives none.

Diagnostic signature. The same name in every minority position. Decisions never change following a dissent. Other members' dissent rate is zero.

Fix. Rotate the steelman role (P-4) so that objection is a function rather than a person. Measure **decision change rate following dissent**; if it is zero over a year, the dissent channel is decorative and the body should stop claiming it.

FM-22.4 — Anonymity Theatre

Description. The first round is nominally anonymous in a group of six, where everyone can identify every author from phrasing and expertise.

Damages. u, plus a false measurement — the group now believes it has an anonymous baseline and draws conclusions from it.

Diagnostic signature. Ask three members to attribute the anonymous responses. **If they are right more than half the time, the protocol is not anonymous.**

Fix. S22.9. Withdraw the anonymity claim, and use P-1's sealed independent estimates instead, which do not depend on anonymity for their effect.

FM-22.5 — Silence Read as Assent

Description. The chair asks "any objections?", hears none, and records consensus. Non-response is tallied as agreement.

Damages. u, `N_eff` — participants who did not speak are counted as if they had.

Diagnostic signature. Minutes record consensus for items where airtime measurement shows several members never spoke. **This is detectable purely from records, retrospectively, and almost nobody checks.**

Fix. S22.8. Record passes as passes. Replace the negative call ("any objections?") with a positive round ("each person, one line").

FM-22.6 — The Correlated Panel

Description. Nine reviewers, all trained in the same three institutions, reading the same journals, funded by the same body. Their agreement is read as convergence on truth.

Damages. u structurally, from the moment of appointment.

Diagnostic signature. Measure error decorrelation, not opinion diversity: give the panel a set of items with known outcomes and check whether their *errors* are independent. **Correlated errors are the definitive reading; correlated opinions are only suggestive.** Then compute `N_eff` and compare with headcount.

Fix. This is a composition problem and no protocol repairs it — the protocols of §22.7 extract the signal present, and if the signals are duplicates there is nothing to extract. Chapter 9's recruitment methods apply: import across training lineages, incentive structures and trades.

FM-22.7 — Protocol Without Storage

Description. Excellent elicitation. Real independent estimates, honest disagreement, a genuine distribution. None of it is written down beyond the final number.

Damages. `eta` — the surplus was generated and not converted.

Diagnostic signature. Ask for last quarter's pre-discussion distributions. If they do not exist, `eta` for this body's most valuable output is zero.

Fix. S22.2 and S22.4. Retain the distributions, retain the reasons, retain the falsifiers. The storage cost is trivial and the audit value compounds.

FM-22.8 — Conscripted Difference

Description. A member is assigned to represent a perspective — demographic, disciplinary, regional — that they did not volunteer to represent and may not hold.

Damages. u (the represented perspective is performed rather than held), and it breaches CANON §5.4.

Diagnostic signature. The member's recorded positions do not differ from the majority except on the items where they are expected to differ.

Fix. Withdraw the assignment. Recruit for genuine difference and let people speak for themselves, which is both the ethical and the measurable option. **Difference that has to be assigned was not present, and assigning it does not create it — it only makes it unmeasurable.**

Worked Example — A Nine-Person Review Panel

A standing panel of nine reviewers assessing proposals against a fixed rubric. Meets monthly. Decisions are consequential and outcomes are eventually observable, which makes the panel measurable in a way most committees are not.

1. Encoding

Choice	Setting
State encoder	Each reviewer's judgement vector across N = 9 rubric dimensions
Interaction observable	Cross-reviewer disagreement per item, measured pre-discussion
Order observable Y	Panel's calibrated accuracy against realised outcomes at 24 months, plus the durable rubric refinements the panel produces
Time scale	One month = one period
Cost observable	Reviewer hours, meeting time, secretariat, appeals

Y is measured from realised outcomes 24 months downstream, entirely independently of the panel's own disagreement record, which satisfies the Paper E §6.3 requirement that the order proxy not be a repackaging of the input.

2. Baseline measurement

Standing parameters, measured over eight periods:

```
N     = 9        ln N = 2.1972
eta   = 0.45   proportion of panel deliberation that becomes durable rubric or precedent
delta = 0.18   turnover of two members a year plus drift in rubric interpretation
C     = 0.42     reviewer time, meeting, secretariat, appeals
Q     = 1.80     accumulated calibrated rubric and precedent
```

u measured by the disagreement-and-error-decorrelation protocol of Chapter 9:

```
u = 0.12
```

Seven of the nine reviewers were trained in three institutions. Six sit on at least one other panel together. Their pre-discussion ratings — recovered from an old pilot exercise — agree within one rubric point on 86 per cent of items.

```
N_eff = 1 + 8(0.12) = 1.96
S     = ln(1.96) = 0.6729
```

Nine reviewers with an effective count of 1.96. The panel is, in the law's accounting, slightly less than two people.

3. The diagnosis

Headroom:

```
h = eta ln N − delta Q − C
  = 0.45(2.1972) − 0.18(1.80) − 0.42
  = 0.9888 − 0.3240 − 0.4200
  = +0.2448
```

Headroom is positive. This is important and counter-intuitive: **the panel is not over-committed. It has budget. It is simply not collecting the input to use it.** Check the step:

```
eta S = 0.45 x 0.6729 = 0.3028
delta Q + C = 0.3240 + 0.4200 = 0.7440
step = 0.3028 − 0.7440 = −0.4412
```

The panel is losing 0.44 units of stored order per period. By Theorem 3.5:

```
Q_inf = (eta S − C) / delta = (0.3028 − 0.4200) / 0.18 = −0.6510
```

A negative equilibrium. Paper E does not hide this failure mode and neither does this manual: **a negative `Q_inf` means the regime is unsustainable as modelled.** The panel's accumulated rubric and precedent will erode to nothing at a rate governed by `tau = 1/delta = 5.6` months, and the erosion will be attributed to turnover, workload, or the declining quality of applications.

None of those is the cause. The cause is that `u = 0.12`.

4. What is required

By Proposition 4.1, achieving an increment `r` per period requires

```
S >= (delta Q + C + r) / eta
```

which is feasible only if `delta Q + C + r <= eta ln N`.

target `r`	required `S`	required `u`	feasible?
0.00 (hold)	1.6533	0.5280	yes
0.05	1.7644	0.6048	yes
0.10	1.8756	0.6906	yes
0.2448 (= `h`, the maximum)	2.1972	1.0000	yes, at orthogonality exactly

Read the first row before anything else. **To merely stop losing ground, this panel needs `u >= 0.528` — it currently has 0.12.** Every month spent below that threshold is a month of erosion, and no amount of reviewer diligence substitutes, because diligence does not appear in the equation.

The last row is the maximum-increment result of Corollary 4.2 arriving as a hard fact: `r_max = h = 0.2448`, and it requires `u = 1`, perfect orthogonality, which is not attainable with real people. **The gauge tells you what you may ask for; the geometry tells you what you can actually deliver; they meet exactly at the ceiling and not before.**

Target: `r = 0.10`, requiring `u >= 0.6906`.

5. The intervention

All five protocols of §22.7, installed together. Total cost increment 0.090, so `C` rises from 0.420 to 0.510.

Protocol	Implementation at this panel
P-1	Each reviewer submits a sealed rating with one line of reasoning 48 hours before the meeting. The distribution is displayed as the first slide of each item.
P-2	The sealed round is unattributed; names attach in round two, optionally.
P-3	Published dissent protection, signed by the commissioning body; dissent is not visible to the appointing authority at reappointment.
P-4	Minority report with a named falsifier on every split decision; consented rotating steelman on unanimous ones.
P-5	Round-robin on each item; the chair speaks last; table replaced with a round one (Chapter 17).

Composition change, per FM-22.6, because protocols cannot extract signal that is not present: two retiring members replaced by a practitioner from an adjacent field and a reviewer from outside the three dominant training lineages.

Measured after four periods: `u = 0.71`.

6. The result

```
N_eff = 1 + 8(0.71) = 6.68
S     = ln(6.68) = 1.8991
eta S = 0.45 x 1.8991 = 0.8546

h = 0.45(2.1972) − 0.18(1.80) − 0.510 = 0.9888 − 0.3240 − 0.5100 = +0.1548
step = eta S − delta Q − C = 0.8546 − 0.3240 − 0.5100 = +0.0206
```

Positive step at the initial `Q`, and it improves as measured `Q` was itself at a depressed level. New steady state:

```
Q_inf = (0.8546 − 0.5100) / 0.18 = 1.9144
```

From a negative equilibrium to `Q_inf = 1.914`. Confirm with the identity of §21.0.2:

```
h at Q_inf = 0.45(2.1972) − 0.18(1.9144) − 0.5100 = +0.1342
eta ln(N / N_eff) = 0.45 x (2.1972 − 1.8991) = 0.45 x 0.2981 = +0.1341   [agrees]
```

Note honestly what happened to headroom: it *fell*, from 0.2448 to 0.1548, because `C` rose by 0.09. **Headroom went down and the system got better.** This is not a contradiction; it is the correct reading of the gauge. Headroom is capacity to build, and the panel spent some of it on the mechanism that lets it build at all. A large unspent headroom in a system with `u = 0.12` is not health. It is a budget nobody can draw on.

7. Accuracy, computed

Take individual reviewer accuracy on the binary "would this proposal have succeeded" question as `p = 0.62`, from the 24-month outcome record. Under full independence, nine reviewers by majority:

```
P_independent(9, 0.62) = 0.7738
```

Under Model C1 with the baseline correlation `rho = 0.75` (badged: a two-point mixture, chosen for transparency, not a general correlation model):

```
P = 0.75(0.62) + 0.25(0.7738) = 0.6584
```

After the intervention, with `rho` estimated at 0.20:

```
P = 0.20(0.62) + 0.80(0.7738) = 0.7430
```

Panel accuracy rises from 65.8 per cent to 74.3 per cent with no change in the individual ability of any reviewer. Seven of the nine are the same people, reading the same proposals, using the same rubric. The change is entirely in the order of operations and the composition of the pool.

Compare against the cascade bound. If the panel deliberated by open sequential discussion with no independent round, §22.4 gives, at `p = 0.62`:

```
P(correct cascade) = [0.62(1.62)/2] / [0.62^2 − 0.62 + 1] = 0.5022 / 0.7644 = 0.6570
```

65.7 per cent — which is, to within a rounding error, exactly the panel's measured baseline of 65.8 per cent. **The panel of nine was performing at the theoretical accuracy of an infinitely long cascade**, which is what a nine-person cascade should do, and is the strongest single confirmation available that the baseline diagnosis was right.

8. Falsification

Declared in advance, per Chapter 25:

1. If measured u does not exceed 0.53 within six periods, the intervention has failed to reach even the hold threshold and must be reported as unsuccessful.
2. If the pre-discussion distributions show the same spread as post-discussion positions did under the old regime, P-1 is not working and Postulate I1's mechanism is wrong in this body.
3. If 24-month outcome accuracy does not rise above the 0.658 baseline within two full outcome cycles, the accuracy model is wrong and should be discarded rather than defended.
4. The shuffled-input control of Chapter 25 is mandatory: if Y is predicted as well by shuffled U as by real U, the entire causal story is a trend and must be reported as such.
5. If the dissent rate rises but decision-change-following-dissent stays at zero, the panel has acquired FM-22.3 and the apparent improvement in u is theatre.

9. The rule, extracted

The panel was not short of ability. Its members were individually accurate at 62 per cent, above the Condorcet threshold, and its headroom was positive throughout. It was failing because nine reviewers had an effective count of 1.96, and because the order of operations in the room guaranteed a cascade before the second item.

Ninety thousandths of a unit of procedural cost, two composition changes, and a round table moved it from a negative equilibrium to a positive one and lifted its accuracy by 8.5 points.

Individuality is not a value the group tolerates. It is the structure the group's function rests on. Merge the individuals and the c-factor falls, the cascade sets in, and the group becomes one mind with one blind spot — and it will feel, from the inside, exactly like agreement.

Chapter 23 . Institutions as Storage

The Law

An institution is not a building, a charter, a payroll, or a mission. In this manual an institution is a **device for holding `Q` against `delta`** — a mechanism whose function is to make stored order survive the departure of the people who made it.

That is the whole definition, and every recognisable institutional form is a special case of it.

Form	What it stores	Against what `delta`
Constitution	Decision procedures and constraints on power	The turnover of every officeholder
Archive	Records of what happened and what was decided	Human memory, at roughly 100% per generation
Curriculum	The transmissible portion of a body of skill	Death of practitioners
Profession	Standards of competence and conduct	Individual variation and market pressure
Standards body	Interface agreements between independent parties	Divergent local optimisation
Maintenance regime	The physical plant and its working order	Wear, corrosion, entropy

Each is `Q` in a different substrate. Each obeys the same recursion, and the parameters differ (CANON §1.4 — this is a claim about model form, not about physics):

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t          S_t = ln(1 + (N−1) U_t)
h_t     = eta ln N − delta Q_t − C_t
```

23.0.1 Why institutions ossify — the equation, not the sermon

Read the headroom gauge with the eyes of an institution that has been successful for sixty years.

```
h_t = eta ln N  −  delta Q_t  −  C_t
      ^^^^^^^^     ^^^^^^^^^     ^^^
      ceiling      maintenance   operating
      fixed        grows with    cost
                   success
```

The first term is a **hard ceiling**. `eta ln N` cannot be raised without limit: `ln N` grows logarithmically, so doubling the number of genuinely distinct categories adds `ln 2 = 0.693` and nothing more, whether you go from 4 to 8 or from 400 to 800 (Chapter 8). `eta` is bounded above by 1 in any sensible normalisation. **The ceiling is essentially fixed for a mature institution.**

The second term is `delta Q_t`, and `Q_t` is the institution's accumulated achievement. Every rule that worked, every precedent that settled a dispute, every procedure that prevented a past disaster, every standard that made two systems interoperate — all of it is `Q`, and all of it requires maintenance at rate `delta` per period, forever.

Therefore: as an institution succeeds, its own stored order consumes an increasing share of a fixed budget. No malice, no decadence, no failure of nerve. Arithmetic.

Ossification is not moral failure. It is a headroom equation.

This is the central claim of the chapter and it is worth stating in its uncomfortable form. The institution that ossifies is being punished *for having worked*. Its `Q` is large because it did its job. The maintenance load is the bill for past success, and it arrives whether or not anyone currently in the building deserves it.

Diagnostically this matters enormously, because the standard response to ossification — replace the leadership, exhort the staff, run a culture programme — addresses a cause that is not operating. **If `delta Q_t + C_t` exceeds `eta ln N`, Corollary 3.4 applies: no arrangement of the people prevents decline.** Not a better chief executive, not a more committed workforce, not a clearer strategy. The constraint is on the budget, and it must be relieved on the budget.

23.0.2 Two identities that structure the whole chapter

Identity 1 — the uniform bound is the zero-headroom line.

Theorem 3.3 gives `Q_t <= max{Q_0, eta ln N / delta}`. Set `C = 0` in the headroom gauge and solve `h = 0`:

```
0 = eta ln N − delta Q     ->    Q = eta ln N / delta
```

The uniform bound on stored order is precisely the level at which headroom reaches zero in a costless system. These are not two facts; they are one fact stated twice. An institution cannot exceed its storage ceiling because at the ceiling it has no budget left to add anything. Every unit of `Q` it holds is a unit of budget it has committed to holding.

For a mature body with `eta ln N = 0.8349` and `delta = 0.015`, the bound is `Q <= 55.66`. That is the absolute maximum stored order the body could hold if it had no operating cost whatsoever. With any real `C`, the practical ceiling is far lower.

Identity 2 — headroom at equilibrium is unused variety.

Derived in §21.0.2 and used throughout here. At any steady state, `delta Q + C = eta S`, so

```
h_inf = eta ln N − eta S = eta ln (N / N_eff)
```

Consequences for institutions specifically:

- A settled institution's headroom equals the difference it is failing to use. Nothing else.
- **A settled institution never has negative headroom.** Therefore, and this is the diagnostic that organises §23.2: **negative headroom is always a disequilibrium condition.** It means `Q` was accumulated under one parameter set and is being carried under a different one. An institution in the red on headroom is always an institution living in a changed world with an inherited balance sheet.

The identity in its general form, and the correction it forces on §23.0.1. Written with the surplus law substituted in, Identity 2 is

```
h_inf = eta ( ln N − F(u*) )              F(u*) = ln(1 + (N−1)u*)
```

and the striking feature is what has vanished. **`delta`, `C` and `Q` all cancel at the fixed point** — verified to six decimal places in Chapter 6, and again here: a body with `eta = 0.35`, `delta = 0.015`, `N = 6`, `u* = 0.55`, `C = 0.42` converges to `Q_inf = 2.840970`, giving `h_inf = 0.164501`, which is `eta(ln N − F(u*))` to the digit.

The consequence corrects the sermon in §23.0.1, and the correction should be read carefully, because the sermon is the one this chapter is named for. `delta Q_t` does grow with success — but under the plain linear law it grows toward a fixed point at which headroom is a *positive* constant that does not depend on `delta`, `C` or `Q` at all. **The plain linear model, left alone, never ossifies.** Maintenance load rises and then stops rising, because `Q` stops rising.

Ossification is therefore not a consequence of accumulation as such. It requires one of two further things, and a diagnosis should say which:

1. **Exogenous `Q` accretion** — obligations added from outside the loop, faster than the loop's own fixed point would settle them: statute, merger, inherited estate, external mandate. `Q` is being pushed past its equilibrium by something that is not `eta S`.
2. **Superlinear cost** — `C` rising faster than linearly in `Q`, which breaks the cancellation. Chapter 40's correction load `C_t = c_0 + kappa Z_t` is one instance; the `n(n−1)/2` coordination term of §23.7 is another.

"Ossification is not moral failure, it is a headroom equation" stands — but the equation names two specific mechanisms, and neither of them is age. Cross-reference Chapter 6, where the identity is derived, and §23.2, whose three routes are exactly the ways a real institution leaves the linear case.

23.0.3 What `h` does not contain

Look at the gauge again and note an absence that decides most institutional strategy:

```
h_t = eta ln N − delta Q_t − C_t
```

`u` does not appear. Non-parallelism enters the storage law only through `S_t = ln(1 + (N−1)U_t)` — the *input* term. It is not in the budget.

The consequence is a hard rule and it contradicts a great deal of received practice:

> **Raising `u` does not raise headroom. It raises the input that fills headroom.**

An institution with `h < 0` that responds by recruiting for diversity has bought a larger tap for a tank that is leaking faster than any tap can fill it. `S` can rise all the way to `ln N` and the step `eta S − delta Q − C` remains negative, because `eta ln N − delta Q − C` — which is exactly `h` — is negative by assumption. **Diversity is the right move at `h > 0` and the wrong move at `h < 0`, and the gauge tells you which you are in before you spend anything.**

The four levers that *do* move `h`, in the order of their usual cost (Chapter 26 develops the ranking):

1. **Raise `eta`.** Usually cheapest. Capture what already happens.
2. **Cut `C`.** Usually available. Remove coordination overhead, friction, unresolved conflict.
3. **Raise `N`.** Expensive and logarithmically rewarded. Rarely the answer.
4. **Shed `Q`.** The one nobody considers, the subject of §23.3, and frequently the only one available at scale.

23.0.4 The fifth lever, and why it is not fifth by accident

The lever missing from that list is the one every institution reaches for first: **lower `delta`.** Build trust, build institutional memory, stop losing what you know. It is absent from the four because it behaves differently from all of them, and the difference is a correction to standard practice rather than a technicality.

Trust does act on `delta` — a body whose members rely on unverified commitments loses less order per period. But `delta` sits in the *denominator* of the steady state, `Q_inf = (eta F(u*) − c*)/delta`, and **when the numerator is negative, dividing by a smaller positive denominator makes the result more negative.** Chapter 49 separates the moves in a worked consortium and computes them:

move	Q_inf
baseline	-0.0836
trust alone (`delta` 0.4300 -> 0.1700)	**-0.2114 — worse**
reciprocity alone (`eta` 0.1721 -> 0.4026)	+0.7394
both	+1.8701

Trust does not add order. It amplifies whichever sign the system already has, and a failing institution that invests in trust first buys a slower, better-mannered decline — together with survey results reporting that the programme worked.

The design rule is therefore an ordering: fix the numerator first. Raise `eta`, or cut `C`, until `eta F(u*) > c*`; only then spend on trust, where the same division now multiplies a positive number. Note also what a lower `delta` does to the diagnosis: it lengthens the time constant `tau = 1/delta`, so the institution takes longer to find out. Cross-reference Chapter 49 §49.14, Chapter 26's "stabilise before growing", and §23.6 below, where the same `delta` that slows the loss also slows the involuntary return of headroom.

The Engineering

23.1 The maintenance term at institutional scale

Institutions are distinguished from teams and rooms by one parameter above all: **their `delta` is small and their `Q` is large.** That is the whole point of them. A constitution with `delta = 0.3` per year is not a constitution.

But `delta Q` is a product, and a small number times a large number is not small.

Take a fixed ceiling `eta ln N = 1.50` and ask what fraction of the annual budget the maintenance term alone consumes:

Q	delta = 0.005	delta = 0.01	delta = 0.02	delta = 0.03	delta = 0.05
10	3%	7%	13%	20%	33%
25	8%	17%	33%	50%	83%

Q	delta = 0.005	delta = 0.01	delta = 0.02	delta = 0.03	delta = 0.05
50	17%	33%	67%	**100%**	**167%**
75	25%	50%	**100%**	**150%**	**250%**
100	33%	**67%**	**133%**	**200%**	**333%**

Every cell at or above 100 per cent is an institution that cannot fund its own maintenance from its own surplus, before paying a single unit of operating cost C. Reading across the delta = 0.02 row: at Q = 10 maintenance is a rounding error; at Q = 50 it is the entire budget.

The same institution, doing nothing differently, moves from the first row to the last row simply by continuing to work. This is the maintenance-vs-growth transition, and every institution that survives long enough reaches it. It is not a phase to be avoided. It is a phase to be *recognised*, because the correct behaviour on either side of it is different and the behaviour that got the institution to the transition is wrong after it.

	Before the transition	After the transition
Dominant term	eta ln N	delta Q_t
Correct posture	Build. Accumulate. Say yes.	Maintain. Curate. Shed.
Success metric	New things established	Old things still working
Typical error	Under-building from caution	Over-building from habit
What "more resources" buys	More Q	A slower decline, briefly
Career incentive	Aligned with the correct posture	**Opposed to it**

The last row is where institutions actually die. **After the transition, the correct organisational behaviour is maintenance and shedding, and no individual career has ever been advanced by maintenance or shedding.** The incentive gradient and the headroom gradient point in opposite directions, and the incentive gradient is the one that people can feel.

23.2 The three routes to negative headroom

Since a steady state cannot have h < 0 (Identity 2), negative headroom always means the parameters changed under an inherited Q. There are exactly three routes, and they are diagnosable apart.

Route A — eta falls Interaction stops converting into durable structure. The meetings still happen; the deliberation is as long as ever; nothing lands.

Mechanisms. Ritualisation — the procedure survives its purpose and is performed rather than used. External override — a higher authority can void local decisions, so deliberation produces nothing durable (this is the failure of Ostrom's Principle 7, Chapter 21). Legal review as a filter that removes anything consequential. Output that is formally produced and never adopted.

Signature. Deliberation hours flat or rising; durable artefacts per hour falling. Ask what last quarter's meetings deposited (Specification S21.7). **An institution whose meetings deposit nothing has eta = 0 and is generating no order at any u whatsoever.**

Route B — `C` rises Operating cost grows, usually as a function of `Q` itself, and usually for defensible reasons: each accumulated rule requires interpretation, each precedent must be checked, each standard must be maintained against its dependents.

Mechanisms. Compliance load proportional to the rule stock. Coordination cost by Brooks's `n(n−1)/2` as the body grows. Unresolved conflict, which compounds (Ostrom P6). Litigation and appeal. The secretariat that grows to service the archive.

Signature. Model `C_t = c_0 + kappa Q_t` and estimate `kappa` by regression across periods. **A significant positive kappa means the institution's own stored order is buying its operating cost**, and it changes the shedding calculation decisively, because shedding `Q` then cuts both the maintenance term and the operating term at once.

Route C — `N` narrows The categorical variety the institution spans contracts. A body founded to arbitrate between fourteen genuinely distinct constituencies now serves nine, because five merged, exited, or were absorbed.

Mechanisms. Industry consolidation. Professional narrowing. Scope creep in reverse — the institution defends a shrinking core. Membership criteria that filter for similarity.

Signature. Count the genuinely distinct categories, not the seats (Chapter 7's redundancy audit). `N` falling from 14 to 9 cuts the ceiling from `eta(2.639)` to `eta(2.197)` — a 17 per cent reduction in the maximum possible input, permanently, and it is almost never noticed because the number of *seats* did not change.

All three routes usually run together, which is why the diagnosis must be quantitative. The Worked Example below shows all three operating simultaneously in one body over thirty-six years.

23.3 Shedding `Q` deliberately

This is the lever nobody reaches for, and at institutional scale it is often the only one with enough range.

The reason it is neglected is not mysterious. `Q` is achievement. Every unit of it was a solved problem, a settled argument, a prevented disaster. Shedding `Q` means deliberately unlearning something that once worked, and the case for keeping each individual unit is always stronger than the case for keeping the aggregate — because the individual case is concrete and the aggregate case is a division.

The arithmetic is unmoved by this. `h_t = eta ln N − delta Q_t − C_t`. If `C_t = c_0 + kappa Q_t`, then

```
dh/dQ = −(delta + kappa)
```

Every unit of `Q` shed returns `delta + kappa` units of headroom, permanently, every period. With `delta = 0.015` and `kappa = 0.0225`, shedding 8 units of `Q` returns `8 x 0.0375 = 0.30` units of headroom per period thereafter. Against a ceiling of 0.83, that is 36 per cent of the entire budget, recovered by deletion.

The shedding instruments **Sunset clauses.** Every rule carries an expiry date. Renewal requires an affirmative act — a positive case, made by someone, at a scheduled time. The mechanism's power is entirely in the default: without a sunset the default is "keep", and keeping requires

no case, no advocate, and no attention. With a sunset the default is "lapse", and now the burden falls on continuation.

A sunset clause does not delete anything. It changes which side of the argument has to do the work, and that is sufficient, because the scarce resource in every institution is not judgement but attention.

Design notes: stagger expiries so review load is level rather than lumpy; make renewal genuinely possible in one step, or the clause becomes a hostage device; and require the renewal case to state what would make the rule unnecessary.

Term limits. Sunset clauses applied to roles rather than rules. Their storage function is double-edged and both edges should be stated: they raise `delta` on role-bound `Q`, which is a cost — experience genuinely leaves — and they prevent role-bound `Q` from becoming permanent, which is the benefit. **A role held for thirty years accumulates a `Q` whose `delta` is zero until the day the holder leaves, and then it is 1.0.** Term limits convert a single catastrophic loss into a series of small scheduled ones, which is precisely the trade a well-designed system makes everywhere else.

Planned demolition. The physical form, and the honest one, because it is visible. A building scheduled for removal at a stated date. A programme with a defined end. A committee constituted for a task and dissolved on delivery. The discipline is that the end date is fixed *at the beginning*, when nobody yet has an interest in extending it.

Forgetting by design. The archive tier. Records are retained but moved out of the default view, so that search cost — a `C` term — does not grow without bound with the archive. This is Chapter 21's evaporation requirement (FM-21.9) at institutional scale: the pheromone trail to the exhausted source must fade, or the colony maintains a motorway to an empty patch of ground.

Deletion budget. The most useful and least used instrument. Fix a rate — *for every new standard adopted, one is withdrawn* — and enforce it structurally rather than by intention. The one-in-one-out constraint keeps `Q` bounded by construction rather than by discipline, which matters because discipline is exactly what fails under schedule pressure.

What must not be shed Shedding is a scalpel and not a policy. Two classes are excluded.

Anything load-bearing for a current dependent. A standard with live implementations is not overhead; it is the interface. Withdrawing it does not shed `Q`, it destroys `Q` elsewhere — usually more than it recovers, and in someone else's balance sheet.

Anything constituting the ability to decline. CANON §5.4 is binding here as everywhere. Procedural protections, appeal rights, dissent records, exit provisions and consent requirements are `Q` that specifically preserves people's capacity to refuse. **They may be simplified; they may not be shed for headroom.** An institution that buys headroom by removing the mechanisms through which people can say no has not improved its parameters. It has changed what it is, and this manual will not recommend it.

The practical form: maintain an explicit **protected set**, written down before any shedding exercise begins, and exclude it from the deletion budget by rule rather than by argument.

23.4 Star & Ruhleder: why the maintenance term is never funded

Susan Leigh Star and Karen Ruhleder, in "Steps Toward an Ecology of Infrastructure" (*Information Systems Research*, 1996), characterised infrastructure relationally rather than as a class of object. Something is infrastructure *for* a given practice, and among its properties: it is embedded, it is transparent in use, it is learned as part of membership, it is linked to conventions, and — the property this chapter turns on — **it becomes visible upon breakdown.**

Mark Weiser's 1991 formulation is the same observation from the other side: the most profound technologies are those that disappear, weaving themselves into the fabric of everyday life until they are indistinguishable from it.

Now put that next to the budget.

```
h_t = eta ln N − delta Q_t − C_t
                 ^^^^^^^^^
                 the maintenance term:
                 large, growing, and invisible while it is being paid
```

The `delta Q` term is the cost of keeping infrastructure working. Infrastructure that is working is invisible. **Therefore the largest and fastest-growing term in an institution's budget is the one with no visible advocate, no constituency, and no photograph.**

The consequences are entirely predictable and they are observed everywhere:

- **Maintenance is deferred first** under budget pressure, because deferring it produces no immediately visible effect. That is not a failure of judgement; it is the correct read of the visible evidence, and the evidence is systematically wrong.
- **Maintenance is measured in avoided events,** and avoided events do not appear in any record. A maintenance regime that works perfectly generates a record indistinguishable from one that was never needed.
- **Maintenance funding is restored only after breakdown,** at which point the cost is the replacement cost rather than the maintenance cost, typically an order of magnitude higher, and the restoration is attributed to the crisis rather than to the neglect.
- **The maintainers cannot describe their contribution** in the vocabulary the institution uses for achievement, because their output is the absence of an event.

The engineering response is to make the invisible term visible *by construction*, since it will never become visible by attention:

Instrument	What it makes visible
A named budget line for `delta Q` — separately reported, not merged into operations	That maintenance has a magnitude
A stated `delta` per asset class, published and revised	That decay has a rate
Scheduled visible failure: a planned outage, a fire drill, a restore-from-archive exercise	That the thing is being held up by something
Deterioration reporting: what got worse this period	That neglect has a direction
Maintenance-to-new-build ratio, reported alongside h	That the two compete for one budget

Scheduled visible failure is the strongest of these and the least used. Star & Ruhleder's finding is that infrastructure becomes visible on breakdown; the engineering inversion is to *cause a small breakdown on purpose, on a schedule*, so that visibility is obtained at a chosen cost rather than at an unchosen one. A quarterly exercise in which a system is restored from the archive tells the institution, in one afternoon, exactly what its `delta` has been doing.

23.5 Goodhart, Campbell, and why measuring institutional health corrupts it

Goodhart's Law: when a measure becomes a target, it ceases to be a good measure. Campbell's Law: the more a quantitative indicator is used for social decision-making, the more subject it is to corruption pressures, and the more apt it is to distort the processes it was meant to monitor.

Institutional health is an unusually soft target, and the corruption is unusually complete, because `Q` is not directly observable. What is observable is a proxy `Y`, and every proxy for stored order has a cheap counterfeit:

Proxy for `Q`	The counterfeit	What it does to the real parameters
Standards published	Publish more, thinner standards	`Q` looks up, `C` genuinely up, `eta` unchanged
Members enrolled	Lower the bar	`N_eff` down (new members duplicate old), `C` up
Procedures documented	Document what nobody does	`Q` flat, `C` up, search cost up
Meetings held	Hold more meetings	`C` up, `eta` down (dilution)
Compliance rate	Narrow the rule until compliance is automatic	`Q` down, measured `Q` up
Archive size	Retain everything	`C` up via search, `delta` masked

Every one of these raises the measured proxy and lowers real headroom. **The counterfeit is not usually fraud. It is the natural drift of an organisation optimising a number in good faith**, and it is fast: Campbell's point is that the corruption is a function of how much weight the indicator carries, not of anyone's intent.

Three defences, in order of strength.

First — measure `h`, not `Q`. Headroom is a difference of three quantities pulling in opposite directions, and it is therefore far harder to counterfeit than any of them. Inflating `Q` *lowers* `h` through the maintenance term. Inflating activity raises `C` and lowers `h`. There is no cheap move that raises measured headroom, because the cheap moves all cost something that `h` subtracts. **A gauge built from a difference resists Goodhart capture in a way a gauge built from a total never can.**

Second — measure `delta` directly, under zero input. Chapter 3's protocol: suspend input to a subsystem, observe the decay of `Y`, and compute `delta = 1 − Y_1/Y_0`. This is expensive and disruptive and it is nearly impossible to game, because it measures what happens when nobody is trying.

Third — rotate the proxy. Announce that Y will be re-specified on a schedule and that the new specification will not be disclosed in advance. This does not prevent optimisation of the current proxy; it prevents the *investment* in optimising it, which is where most of the distortion lives.

And a standing honesty requirement, per CANON §5.3: publish the proxy, publish its known counterfeits, and publish the falsification rules from Chapter 25 alongside every reported figure. **A framework that cannot fail cannot be trusted.**

23.6 The two ways headroom returns

Here is the structural fact that organises the Worked Example, and it is the chapter's thesis.

An institution in negative headroom does not stay there forever. The storage law will not permit it. If `h_t < 0` then `eta S_t < delta Q_t + C_t` for every attainable `S_t`, so `Q_{t+1} < Q_t`: stored order falls. And as `Q` falls, `delta Q` falls, and `h` rises. Eventually `h` crosses back above zero.

Headroom always returns. The only question is whether the `Q` that was spent to buy it back was chosen or surrendered.

	Involuntary	Deliberate
Mechanism	Decay: the surplus does not cover maintenance, so things quietly stop working	Shedding: rules sunset, programmes end, roles rotate
Timing	Set by `delta` and by the size of the gap	Chosen
What is lost	Whatever was weakest, least defended, or least visible	Whatever was chosen, by stated criteria
Experienced as	Decline, drift, irrelevance, "we're not what we were"	A decision, with an argument attached
Duration in `h < 0`	Years to decades	Zero
Protected set	Not protected — decay does not read the exclusion list	Protected by construction
Attribution	To the current leadership, who did not cause it	To the decision, correctly

The third row is the one to hold on to. **Involuntary shedding does not shed the least valuable `Q`.** It sheds whatever had the fewest defenders — which correlates with invisibility, which by Star & Ruhleder correlates with *working well*. The archive that never failed loses its funding before the committee that meets monthly and produces nothing, because the committee has members and the archive has none.

Deliberate shedding is the same arrival, chosen, earlier, and with control over what is lost. That is the entire case for sunset clauses, term limits, and deletion budgets, and it does not require anyone to believe that the shed rules were bad. It requires only that the budget is finite and that someone should choose.

23.7 The institutional survey, in brief

Full protocol in Chapter 24; the institution-specific points are these.

`Q` proxy. For a rule-holding institution the natural proxy is the count of live rules weighted by dependency. **Weighting is not optional:** an unweighted count treats a foundational standard with two hundred implementations as equal to a procedural note nobody has read, and it will therefore recommend shedding the wrong things.

Justify that `Y` is not the input restated (Paper E §6.3). For a standards body, `U` is the disagreement structure among member organisations and `Y` is the live weighted rule stock. These are plainly different observables — a year of intense disagreement producing no adopted standard scores high `U` and zero `Delta Y`, which is the required test.

`delta` estimation. Institutions rarely permit a zero-input period, so use natural experiments: a dormant working group; a subsystem that lost its staff; the period after a merger when a function was unattended. **`delta = 1 − (Y_1/Y_0)` over the dormant interval, annualised.**

`C` tally. Include the items nobody budgets: unresolved dispute load, search cost across the archive, compliance interpretation, and the coordination cost `n(n−1)/2` of the governing body itself. These are typically the majority of real `C` and typically absent from the accounts.

`N` audit. Count genuinely distinct constituencies, not seats. Run Chapter 7's redundancy audit and expect the honest `N` to be materially below the headcount.

The cheap instrument, and it is cheaper than everything above. Chapter 42 derives the headroom identity `grad h = −grad C`: of the three terms in the gauge, only `C` depends on configuration, so the gradient of headroom is the negative of the gradient of cost. **Therefore a habit survey is a headroom measurement.** Where do people actually route their requests; which form gets filled and which gets worked around; which meeting is attended and which is dialled into from a corridor; where has the paving been worn through the grass. Those recurring choices are the institution's desire lines, and they point down the cost field — which is to say, up the headroom field.

The practical consequence is that **the full survey of Chapter 24 is not always required to locate the direction of steepest improvement**, only to price it. A tally of routed-around procedures, costing an afternoon, identifies the same direction that an instrumented survey identifies at far greater expense. Two conditions, both stated plainly in Chapter 42: the identity holds only where `eta`, `N`, `delta` and `Q` do not themselves move with the configuration — a workaround that destroys a distinction moves `N`, and the identity fails there — and the tally recovers cost *differences*, not absolute costs, so it ranks candidate moves and does not budget them. Cross-reference Chapter 42 §42.0.6 and Chapter 24.

Reporting. Report `h` with its three components separated, every period, with the maintenance term shown as a percentage of the ceiling. **That single percentage predicts institutional behaviour better than any strategy document.**

The Specification

S23.1 — Report headroom, decomposed

The institution **shall** report, each period: `eta ln N`, `delta Q_t`, `C_t`, and `h_t`. The maintenance term **shall** be reported as a percentage of the ceiling. **Reporting `Q` alone is prohibited under this specification**, because `Q` alone is the Goodhart-capturable quantity.

S23.2 — Named maintenance line

`delta Q_t` **shall** appear as a distinct, named budget line, not merged into operations, and **shall** be defended separately in every budget cycle.

S23.3 — Published decay rates

For each asset class holding `Q`, the institution **shall** publish an estimated `delta` and the method used to estimate it, and **shall** revise it at least every three years.

S23.4 — Sunset by default

Every new rule, standard, procedure or programme **shall** carry an expiry date at adoption. Renewal **shall** require an affirmative case, made at a scheduled time by a named advocate. Expiries **shall** be staggered so that review load is level.

S23.5 — Deletion budget

A withdrawal rate **shall** be fixed and enforced structurally. The default requirement is **one-in-one-out by dependency-weighted volume**. The rate may be varied only by an explicit decision recorded with its headroom justification.

S23.6 — Protected set

Before any shedding exercise, a written protected set **shall** be produced, containing at minimum: (a) rules with live external dependents; (b) all procedural provisions that constitute a person's ability to decline, appeal, dissent or exit. **The protected set is excluded from the deletion budget by rule and not by argument** (CANON §5.4).

S23.7 — Term limits with handover deposit

Where term limits apply, the outgoing holder **shall** deposit role-bound `Q` in the shared medium (Chapter 21) before departure, and the handover **shall** be verified by the successor's ability to reconstruct the state without the predecessor present. **A term limit without a handover deposit converts a slow leak into a scheduled catastrophe.**

S23.8 — Scheduled visible failure

At least annually, the institution **shall** conduct an exercise in which a piece of working infrastructure is deliberately made visible: a restore from archive, a planned outage, a reconstruct-from-record drill. Findings **shall** be reported against the published `delta` estimates of S23.3.

S23.9 — The h < 0 rule

When measured h_t < 0, the institution **shall not** commission new durable structure. Permitted actions are exactly: raise eta, cut C, shed Q. **Recruiting for diversity is explicitly not a permitted response to negative headroom** (§23.0.3), because u does not appear in h.

S23.10 — Proxy honesty

Every reported proxy **shall** be published together with its known counterfeits (§23.5) and with the falsification rules of Chapter 25, including the shuffled-input control.

The Failure Modes

FM-23.1 — The Maintenance Trap

Description. delta Q_t has grown to consume the entire budget. The institution maintains and cannot build. Every proposal is met with "we don't have capacity," which is precisely correct and is mistaken for a cultural problem.

Damages. h (to zero or below), and then Q involuntarily.

Diagnostic signature. delta Q_t / (eta ln N) >= 0.9. Corroborating symptoms: no new durable output for several periods despite unchanged staffing; strategy documents that repeat across years; a rising share of effort in interpretation and support of existing structure.

Fix. Shed Q (S23.5) and cut C. **Do not attempt to raise the ceiling** — ln N is logarithmic and will not deliver, as Chapter 8 computes. This is the failure mode for which deliberate shedding was invented.

FM-23.2 — The Unshed Q

Description. Rules, programmes and roles accumulate without any withdrawal mechanism. Nothing has been removed in living memory. The rule stock is a stratigraphy in which nobody can tell live from fossil.

Damages. delta Q (rises without bound to the uniform limit), C (search and interpretation cost), eta (new work must be reconciled against everything).

Diagnostic signature. Count withdrawals over the last five years. **Zero is the fault.** Second test: sample ten rules at random and count how many have a live dependent; a hit rate below half means most of the stock is fossil.

Fix. S23.4 and S23.5 — sunset by default and a deletion budget. Begin with the sample: withdrawing ten rules that have no dependents is uncontroversial and establishes that withdrawal is possible, which is the actual barrier.

FM-23.3 — Invisible Infrastructure

Description. Maintenance is unfunded because working infrastructure is invisible (Star & Ruhleder). The budget is cut, nothing happens for two years, the cut is judged a success, and the cut is

repeated. Then something fails completely.

Damages. `delta` (effectively rises — unmaintained order decays faster), `Q` (step loss at breakdown), `C` (replacement at crisis prices).

Diagnostic signature. Maintenance budget falling in real terms across periods with no measured rise in `delta` — because nobody is measuring `delta`. **The absence of the measurement is the signature.**

Fix. S23.2, S23.3 and S23.8. The scheduled visible failure is the one that changes minds, because it converts an argument into an observation.

FM-23.4 — Goodhart Capture

Description. A proxy for institutional health became a target. The proxy now rises reliably while the institution declines.

Damages. All parameters, and the measurement system that would have detected the damage.

Diagnostic signature. The reported metric improves monotonically across periods — real institutional quantities are noisy, and **a smooth monotone series is evidence of management of the series, not of the thing.** Cross-check: does an independent outcome measure move with the proxy?

Fix. S23.1 — report h rather than any total. Rotate the proxy on an undisclosed schedule. Publish the counterfeits.

FM-23.5 — Ossification Diagnosed as Culture

Description. An institution with `h < 0` is treated to a leadership change, a values programme, a restructure, or an exhortation to be more innovative. None of it works. The next leadership is blamed and the cycle repeats.

Damages. `C` (the interventions themselves are expensive), morale, and time — the delay while the real constraint continues to bind.

Diagnostic signature. Compute `h`. If `delta Q_t + C_t > eta ln N`, **Corollary 3.4 applies and no arrangement of the people prevents decline.** The signature is that successive leaderships of differing style produce indistinguishable results, which is exactly what a budget constraint predicts and a culture theory does not.

Fix. Relieve the budget: shed `Q`, cut `C`, raise `eta`. Then, and only then, is a leadership or composition question worth asking.

FM-23.6 — Diversity Recruited Into a Deficit

Description. An institution with negative headroom responds by broadening its membership and recruiting for difference. `u` rises. Nothing improves. The failure is then attributed to the new members or to resistance from the old.

Damages. `C` (a larger body costs quadratically more to coordinate), and the credibility of u-raising interventions for a generation.

Diagnostic signature. `h < 0` before the intervention; `u` measurably higher after; `Q` still falling.

Fix. Sequence correctly, per S23.9 and Chapter 26: **stabilise before growing.** Bring `h` above zero by shedding and by `eta` and `C` work, and *then* raise `u` to fill the budget you have created. The diversity intervention is not wrong; it was administered in the wrong order.

FM-23.7 — The Sunset That Never Sets

Description. Sunset clauses exist and every renewal is automatic. The review is a formality, the advocate is the secretariat, and no rule has ever lapsed.

Damages. `C` (the review process is pure overhead), plus a false sense that FM-23.2 has been addressed.

Diagnostic signature. Renewal rate at or near 100 per cent. **A functioning sunset regime lapses things; a renewal rate above roughly 90 per cent means the mechanism is decorative.**

Fix. Require the renewal case to be made by a named beneficiary who is not the secretariat, and require it to state a live dependent. Pair with S23.5, which forces a withdrawal whenever an adoption occurs and therefore cannot be satisfied by universal renewal.

FM-23.8 — Term Limits Without Deposit

Description. Rotation is enforced. Nothing was written down. Each rotation loses the role's accumulated judgement entirely.

Damages. `Q` (step loss per rotation), `delta` (effectively raised to the rotation rate on all role-bound order).

Diagnostic signature. Performance sawtooth synchronised with the rotation cycle. Bus-count of 1 on role-bound state (Chapter 21, S21.6).

Fix. S23.7. Term limits and shared media are complements, not alternatives: **rotation without deposit is not a shedding strategy, it is a leak with a calendar.**

FM-23.9 — The Protected Set Raided

Description. Under headroom pressure, the shedding exercise reaches appeal rights, dissent procedures, consultation requirements and exit provisions, because they are procedurally expensive and have no organised defenders.

Damages. Breaches CANON §5.4. Also, in the model's own terms: it destroys the u-preserving structure of Chapter 22 and therefore the input term, so it does not even deliver what it promises.

Diagnostic signature. The shedding list contains items whose function is to let someone say no.

Fix. S23.6, enforced structurally: the protected set is written before the exercise begins and is excluded by rule. **Headroom bought by removing people's ability to decline is out of scope for this manual, and no arithmetic in it licenses that trade.**

Worked Example — A Technical Standards Body, 1962 to 2026

A standards body founded in 1962 to publish interface standards for a manufacturing sector. Sixty-four years of continuous operation. All figures below are computed; the parameter values are a constructed case rather than a measured organisation, and are stated as such.

1. Encoding

Choice	Setting
State encoder	Member organisations grouped into distinct constituencies by manufacturing role
Interaction observable	Cross-constituency working-group participation and recorded technical disagreement
Order observable Y	Live standards weighted by number of external implementations
Time scale	One year = one period
Cost observable	Secretariat, working-group time, interpretation requests, appeals — normalised to Y units

Y (dependency-weighted live standards) is independent of U (disagreement structure among members), satisfying Paper E §6.3: a year of intense committee disagreement producing no adopted standard registers high U and zero `Delta Y`.

2. Era 1 — 1962 to 1990

Founding parameters, estimated from the archive:

```
N     = 14        fourteen genuinely distinct constituencies    ln N = 2.6391
u     = 0.60       constituencies with materially different interests
eta   = 0.62       working groups produced adopted standards at high rate
delta = 0.030      standards decay slowly; technology drift and disuse
C     = 0.320      small secretariat, low interpretation load
Q_0   = 0          nothing stored at founding
```

```
N_eff = 1 + 13(0.60) = 8.80
S     = ln(8.80)     = 2.1748
eta ln N = 0.62 x 2.6391 = 1.6362      <- the ceiling
eta S    = 0.62 x 2.1748 = 1.3483      <- the input
```

Trajectory:

Year	Q_t	h_t
1962	0.000	+1.3162
1967	4.842	+1.1709
1972	9.001	+1.0462
1976	11.900	+0.9592

Year	Q_t	h_t
1982	15.638	+0.8471
1990	**19.669**	**+0.7261**

Twenty-eight years of accumulation. Headroom positive throughout and falling steadily — **success consuming its own headroom, exactly as Chapter 6 describes.** Steady state for these parameters would be `Q_inf = (1.3483 − 0.320)/0.030 = 34.28`, with `h_inf = eta(ln N − S) = 0.62(0.4643) = 0.2879` by Identity 2. The body never reached it.

Maintenance share at 1990: `delta Q / (eta ln N) = 0.030 x 19.669 / 1.6362 = 36%`. Substantial and not yet critical.

3. Era 2 — the parameters change, the `Q` does not

Between roughly 1985 and 1992 the sector consolidated, and three of the chapter's named routes operate at once.

Route C — `N` narrows. Fourteen constituencies become nine through merger and exit. The ceiling falls from `eta(2.6391)` to `eta(2.1972)`.

Route A — `eta` falls, 0.62 to 0.38. Standards now require legal review, patent clearance and a formal objection period. Working-group output that once became a standard in eighteen months now takes four years, and a growing share never completes.

Route B — `C` becomes a function of `Q`. With 19.7 units of live dependency-weighted standards, interpretation requests, conformance disputes and revision maintenance now scale with the stock:

```
C_t = c_0 + kappa Q_t,     c_0 = 0.320,   kappa = 0.0225
```

And `delta` is deliberately driven down, 0.030 to 0.015, by a policy that no standard is ever withdrawn — which the body regards, reasonably, as a service to its implementers.

Also u falls from 0.60 to 0.28: the surviving members are large integrated firms with substantially aligned interests.

```
N     = 9          ln N = 2.1972
u     = 0.28  ->  N_eff = 1 + 8(0.28) = 3.24,   S = ln(3.24) = 1.1756
eta   = 0.38  ->  eta ln N = 0.8349,   eta S = 0.4467
delta = 0.015
C_1990 = 0.320 + 0.0225(19.669) = 0.7626
```

Headroom at 1990:

```
h = 0.8349 − 0.015(19.669) − 0.7626
  = 0.8349 − 0.2950 − 0.7626
  = −0.2226
```

Negative. The body did nothing wrong in 1990. It inherited a `Q` of 19.7 accumulated under a ceiling of 1.6362 and is now carrying it under a ceiling of 0.8349. That is Identity 2's corollary in the flesh: **negative headroom is always inherited order under changed parameters.**

The zero-headroom stock level for the new parameters:

```
Q* = (eta ln N − c_0) / (delta + kappa) = (0.8349 − 0.320) / 0.0375 = 13.73
```

The body holds 19.67. **It is 5.94 units of stored order over the line that its current parameters can support**, and every one of those units was a solved problem.

4. Path A — do nothing

The body continues. Standards are neither withdrawn nor renewed; working groups meet; adoption is slow. `Q` decays because `eta S = 0.4467` does not cover `delta Q + C`.

Year	Q_t	C_t	h_t	Maintenance + operating as % of ceiling
1990	19.669	0.7626	**-0.2226**	127%
1995	16.835	0.6988	**-0.1164**	114%
2000	14.495	0.6461	**-0.0286**	103%
2005	12.561	0.6026	+0.0439	95%
2007	11.885	0.5874	+0.0692	92%
2010	10.964	0.5667	+0.1038	88%
2015	9.644	0.5370	+0.1533	82%
2020	8.554	0.5125	+0.1942	77%
2026	**7.494**	0.4886	+0.2339	72%

Twelve consecutive years of negative headroom, 1990 through 2001. Twelve years in which the body could not, by Corollary 3.4, commission any durable new structure at any interaction geometry whatsoever. During those twelve years the body appointed three directors-general, ran two strategic reviews and one culture programme. This is FM-23.5 in its natural habitat.

Headroom did return, in 2002, and by 2026 it stands at a healthy +0.2339. Read the price:

```
Q: 19.669 (1990)  ->  7.494 (2026)
Lost: 12.175 units = 61.9% of everything accumulated in twenty-eight years
```

The body bought back its headroom by losing 62 per cent of its stored order, over thirty-six years, without choosing a single unit of what was lost. What lapsed was whatever had the fewest defenders — which, per Star & Ruhleder, correlates with whatever was working quietly.

5. Path B — deliberate sunset in 1990

Counterfactual. In 1990 the body computes h, finds `−0.2226`, and applies S23.4 through S23.6: a protected set is written (standards with live implementations, all appeal and objection procedures), and 40 per cent of dependency-weighted stock is sunset in a single scheduled exercise.

```
Q: 19.669 -> 11.801
C: 0.320 + 0.0225(11.801) = 0.5855
h = 0.8349 − 0.015(11.801) − 0.5855 = +0.0724
```

Positive immediately. Trajectory:

Year	Q_t	h_t
1990	11.801	+0.0724
1995	10.336	+0.1273
2000	9.126	+0.1727
2010	7.301	+0.2412
2026	5.507	+0.3084

Zero years of negative headroom. But note the honest result, and it is the reason this path is not the recommendation: **Path B ends in 2026 with `Q = 5.507`, which is *less* than Path A's 7.494.**

Shedding alone bought thirty-six years of positive headroom and then spent none of it. `eta` was still 0.38, u was still 0.28, so `eta S` was still 0.4467 and still insufficient. The body was solvent and idle.

> **Shedding buys headroom. Headroom unspent is headroom wasted. Shedding is not a strategy; it is the precondition for one.**

6. Path C — shed, then spend

The full intervention, sequenced per S23.9 and Chapter 26: stabilise first, then grow.

1. **Shed 40 per cent** of dependency-weighted stock, protected set excluded. `Q: 19.669 -> 11.801`.
2. **Raise `eta` from 0.38 to 0.60.** Working-group output is captured directly into draft standards with a standing editor; the objection period runs concurrently with drafting rather than after it; every working group deposits to the shared medium (Chapter 21, S21.7). This is `eta` work — the cheapest lever — and it does not require any change in what anyone is trying to do.
3. **Raise u from 0.28 to 0.55** — *only now*, with h positive. Membership is opened to integrators, testing laboratories and two adjacent-sector bodies. Chapter 22's protocols run in every working group.

```
eta ln N = 0.60 x 2.1972 = 1.3183
N_eff    = 1 + 8(0.55)   = 5.40
S        = ln(5.40)      = 1.6864
eta S    = 0.60 x 1.6864 = 1.0118
h_1990   = 1.3183 - 0.015(11.801) - 0.5855 = +0.5558
```

Year	Q_t	C_t	h_t
1990	11.801	0.5855	+0.5558
1995	12.958	0.6116	+0.5124
2000	13.913	0.6330	+0.4766
2010	15.354	0.6655	+0.4226
2026	**16.770**	0.6973	**+0.3695**

Steady state, by Theorem 3.5:

```
Q_inf = (eta S - c_0) / (delta + kappa) = (1.0118 - 0.320) / 0.0375 = 18.449
```

Identity 2 check:

`h_inf = eta ln(N / N_eff) = 0.60 x (2.1972 − 1.6864) = 0.60 x 0.5108 = +0.3065`

The trajectory is converging on `Q_inf = 18.449` with `h_inf = +0.3065`, consistent with the table. Time constant `tau = 1/(delta + kappa) = 26.7` years — institutions are slow, and the slowness is the same property that makes them worth having.

7. The three paths compared at 2026

	Path A: do nothing	Path B: shed only	Path C: shed then spend
Years with `h < 0`	**12**	0	0
`Q` in 2026	7.494	5.507	**16.770**
`h` in 2026	+0.2339	+0.3084	+0.3695
`Q` lost from 1990 peak	61.9%	72.0%	**14.7%**
`Q` in 2026 vs post-shed 1990 base	—	-53.3%	**+42.1%**
What was lost, chosen?	No	Yes	Yes
Protected set honoured?	No	Yes	Yes

Path A lost more than Path C and chose none of it. Path B was disciplined and inert. **Path C ends with 2.2 times Path A's stored order and higher headroom, having deleted 40 per cent of its rule stock on purpose in a single year.**

The lesson is not that shedding is good. It is that **the budget is finite, headroom returns either way, and the only decision available is whether the institution chooses what it gives up or discovers it thirty years later.**

8. What would falsify this

Declared in advance, per Chapter 25:

1. If the dependency-weighted stock `Y` does not decay at approximately `delta = 0.015` per year during a period of suspended adoption, the `delta` estimate is wrong and every projection above is void.
2. If `C_t` does not regress significantly on `Q_t` — that is, if `kappa` is not distinguishable from zero — then Route B is not operating in this body, and the shedding case loses most of its arithmetic force. **kappa must be estimated, not assumed.**
3. If, after the `eta` intervention, adopted standards per working-group hour does not rise, `eta` did not move and the intervention failed regardless of what the process documents say.
4. If measured `u` rises after the membership change but `Y` does not, then either `eta` is lower than estimated or the new members are duplicating existing constituencies — run Chapter 7's redundancy audit before concluding either.
5. The shuffled-input control (Chapter 25) is mandatory: if `Y_{t+1}` is predicted as well from a shuffled `U` series as from the real one, the model is fitting a trend and the entire analysis must be reported as unsupported.

9. The rule, extracted

The standards body of 1990 was not badly run. It was a successful institution carrying the stored order of a world that had changed. Its ceiling had fallen by 49 per cent — from 1.6362 to 0.8349 — and its maintenance and operating load had not.

Nobody in the building could see it, because Q is invisible while it works, and the maintenance term has no advocate. What they could see was that everything took longer, that proposals stalled, and that morale was poor. They diagnosed culture. They were looking at arithmetic.

Ossification is not moral failure. It is a headroom equation. Measure h, decompose it, and if it is negative, do not build — shed, and choose what you shed, because the alternative is not keeping it. The alternative is losing it slowly to whatever had the fewest friends.

Part V — Tradecraft

Chapter 24 . The Survey

You do not get to touch the system until you have measured it.

That is not a moral position. It is an arithmetic one. Every design move available to you — raising conversion, cutting cost, widening the category set, shedding stored order — changes a term in an inequality whose sign you cannot guess. **A system with negative headroom and a system with generous headroom look identical from the outside; they differ only in numbers you have not yet taken.** The survey is the procedure that takes them.

This chapter is the first of the three tradecraft chapters. It tells you how to walk into a real organisation, a real room, or a real institution and come out with six numbers — delta, eta, N, u, C, Q — and one verdict, h. Chapter 25 tells you how to commit to a model of those numbers in a way that can fail. Chapter 26 tells you what you are permitted to build once you know.

The Law

The equation you are surveying

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t      S_t = F(U_t) = ln(1 + (N−1) U_t)

Q_t      stored order            delta   retention loss per period, in [0,1]
S_t      surplus input            eta     conversion efficiency, >= 0
C_t      cost / dissipation       N       number of categories
U_t      interaction geometry in [0,1],  S_t in [0, ln N]
```

This is **Paper E §3**, and it is proven mathematics under its own assumptions. The exact solution (Theorem 3.1), the one-step growth condition (Corollary 3.2), the uniform bound (Theorem 3.3), the sustainability ceiling (Corollary 3.4) and the steady state (Theorem 3.5) are all theorems. They are theorems *about the equation*. They become statements about your organisation only if the equation is a faithful model of it, and **whether it is a faithful model of it is exactly what the survey and the preregistration exist to decide.**

The master gauge (Paper E §6.4):

```
h_t := eta ln N − delta Q_t − C_t

h_t <  0    unsustainable in this model. No interaction geometry saves it.
h_t =  0    maintenance only. Hold; do not build.
h_t >  0   growth available; any requested increment must satisfy r_t <= h_t.
```

The five required operational choices (Paper E §5.1)

Paper E does not permit you to apply the storage law to a real system by analogy. Before any number is taken, five choices must be written down, in advance, in language specific enough that a second engineer could take the same measurements and get the same numbers. They are:

#	Choice	What it fixes	Symbol produced
1	**State encoder**	how each participant, discipline, or category becomes a vector	`x_i`
2	**Interaction observable**	how the encoded states become a single geometry number per period	`U_t`
3	**Order observable**	what durable thing you will count as stored order	`Y_t` (proxy for `Q_t`)
4	**Time scale**	how long one period is	`t`
5	**Cost observable or cost model**	how dissipation is counted, in the same units as `Y`	`C_hat_t`

These five choices are the entire empirical content of the framework. The mathematics is fixed. Everything that can be wrong about your application of it is in these five lines. An engineer who writes them down badly has not produced a bad estimate; they have produced no estimate at all.

The proxy postulate — badged

POSTULATE (badged, house rule from Book One: the modelling assumption is permitted, but it is permitted in the open, with its badge on). The survey assumes that the stored order of a real system — a body of practice, a trained staff, a furnished room, a working institution — is adequately summarised by a single non-negative scalar `Y_t`, and that this scalar moves according to the storage law up to noise. **Nothing in Papers A-E proves this.** Paper A is pure mathematics and makes no empirical claim. Paper E proves properties of a recursion; it does not prove that any organisation obeys it. The scalar-adequacy assumption is a design decision you make, and Chapter 25 is the procedure by which you allow reality to reject it.

The honesty clauses, restated because the survey is where they bite

- The triangle inequality **fails** for geometric triples (Paper A §5). Subadditivity holds on the parameter `u` only. Do not compose surplus across three groups by adding distances.
- As deployed, the surplus functional is a **lower bound** on the ensemble's information, not the information itself. **We log the floor, not the ceiling.** A survey that reports `F(u)` is reporting a guaranteed minimum, and the true richness of the interaction may be higher.
- Paper D: surplus is **not** identified with Shannon entropy or mutual information in general. Do not import entropy intuitions into your cost accounting.
- Paper E proves **no thermodynamic negentropy law**. `F(u) + N_can(u) <= ln N` (Corollary 2.2); higher static surplus does not imply higher canonical concentration.
- `0 <= F(u) <= ln N` (Paper D, Corollary 2.2). **ln N is the hard ceiling on surplus per interaction**, and it is why `h` is defined with `ln N` rather than with `F(U_t)`.

The one law-level warning: Paper E §6.3

The model is only a claim about the world if `Y` is not a repackaging of `U`. If your order observable is computed from the same records, by the same instrument, at the same moment as your interaction observable, then `Y_{t+1} = a Y_t + b F(U_t) − C_t` can fit beautifully and mean nothing at all. It will have measured an algebraic relationship between two encodings of one measurement. §24.15 gives the four tests that prevent this, and the Worked Example of Chapter 25 shows a repackaged proxy passing every statistical baseline in the book.

The Engineering

24.1 What a survey is, and what it is not

A survey is not an audit, an assessment, a maturity model, or a culture questionnaire. It produces six numbers and one verdict. It has a fixed output format. It is finished when the format is full and every entry has a stated provenance.

It is also not a diagnosis. The survey tells you the state of the accumulator. It does not tell you whether the accumulator should be full. **An institution with a low `Q` may be young, or may be hollowed out, and the survey cannot tell you which — only the trajectory can.** That is why the survey requires a *series*, not a snapshot.

Minimum survey deliverable:

```
SURVEY OUTPUT (one page)
-------------------------------------------------------------------
System ............. name, boundary, and what is excluded
Period ............. one period = ______ ; series length ______
Encoder ............ how x_i was formed
N .................. ______   ln N = ______
U (capacity) ....... ______   from ______ instrument
U (realised) ....... ______   from ______ instrument
F(U) realised ...... ______   = ln(1 + (N-1) U)
delta .............. ______ /period   half-life ______ periods
eta ................ ______ Y-units per nat
C_hat .............. ______ Y-units/period  (denominated in Y)
Q (= Y at survey) ... ______ Y-units
-------------------------------------------------------------------
eta ln N  = ______
delta Q   = ______
h         = ______        VERDICT: [ h<0 | h=0 | h>0 ]
one-step  = eta F(U) - delta Q - C_hat = ______
Q_inf     = (eta F(U) - C_hat)/delta   = ______
-------------------------------------------------------------------
```

Everything in this chapter exists to fill that page defensibly.

24.2 Choice 1 — the state encoder

The state encoder turns each *category* into a unit vector. It is the choice that determines what "different" means for your system, and it is therefore the choice with the largest downstream effect on u.

Three encoders are practical in the field.

Encoder A — the disposition vector (recommended default). Assemble a fixed probe set of m real cases the system must handle. Ask each category (each discipline, each role, each site) to record its disposition on each case as +1 or −1. The category's vector is that list of m signs, divided by sqrt(m) so it is a unit vector. Then for two categories i and j that agree on m − k cases and oppose on k:

```
x_i . x_j = (agreements − disagreements) / m = 1 − 2 d_ij     where d_ij = k/m

u_ij = 1 − (x_i . x_j)^2
```

This encoder is exact, not approximate: for a plus/minus encoding the inner product *is* one minus twice the disagreement rate.

Encoder B — the error vector. For systems that produce forecasts or judgements with later ground truth, encode each category by its vector of signed errors over the last m resolved cases, mean-centred and normalised. This is the strongest encoder available, because decorrelated error is the operational meaning of independent judgement (Condorcet's independence condition, Hong and Page's diversity result, and Woolley's collective-intelligence factor all point at the same quantity from different directions).

Encoder C — the practice vector. For systems with no shared probe set, encode by observable practice: which tools, which sources, which sequence of steps, which vocabulary. Score each category on m binary practice items. Weakest of the three, because two people can share every practice and still disagree — and can share no practice and still agree.

Encoder warning that catches most beginners. u = 1 − cos^2 is the squared *sine* of the angle between the encoded categories, and the sine rises steeply near alignment. Under Encoder A, a disagreement rate of only 1 case in 12 already gives

```
d = 0.0833 -> cos = 0.8333 -> u = 1 − 0.6944 = 0.3056
```

and a disagreement rate of 3 in 12 gives u = 0.75. **The encoder is generous: modest observed disagreement produces large u.** This is not an error in the mathematics; it is a property of the geometry, and it means two things for you. First, you must declare your encoder before you measure, because a different encoder gives a different u for the same organisation. Second, a low measured u under Encoder A is a very serious finding: it means the categories are near duplicates of each other, not merely similar.

24.3 Choice 2 — the interaction observable

You now have N unit vectors and you need one scalar U_t in [0,1] per period. The framework requires a scalar. It does not tell you which aggregator to use, so **the aggregator is a declared modelling choice, badged as such, and it must appear in the preregistration.**

The recommended aggregator is the mean pairwise non-parallelism over all ordered pairs actually present in the period:

```
U_t = mean over i != j of ( 1 − (x_i . x_j)^2 )
```

There is a good reason to prefer it beyond convenience. Let `G` be the Gram matrix of the `K` unit vectors, so `G_ij = x_i . x_j` and `tr G = K`. The participation ratio of the Gram spectrum — the standard Renyi-2 effective count — is

```
PR = (tr G)^2 / ||G||_F^2 = K^2 / ||G||_F^2
```

Since `||G||_F^2 = K + sum_{i != j} G_ij^2 = K + K(K−1)(1 − U)`, a line of algebra gives

```
K / PR = 1 + (K−1)(1 − U)        hence        U = 1 − (K/PR − 1)/(K−1)
```

and therefore, for the effective count of the Surplus Law,

```
N_eff = 1 + (N−1) U = N + 1 − N/PR
```

The mean pairwise aggregator is not an arbitrary average. It is exactly the participation-ratio (Renyi-2) diversity of the encoded ensemble, reparametrised into the effective-count form the Surplus Law uses. This identity is exact in double precision; it was verified numerically for `K = 3, 5, 8, 12` on random ensembles with agreement to better than 1e-12. Declaring the mean pairwise aggregator therefore commits you to a standard and well-understood diversity measure rather than to an improvisation.

Record it as a named identity, because Chapter 7 rests on it. Chapter 7 defines the effective count as the roster audited for redundancy; `N_eff = N + 1 − N/PR` says exactly what that audit is, in a form an ecologist, a portfolio analyst or a spectroscopist would recognise on sight. Two limiting cases follow and both belong on the survey form. At `PR = N` — a Gram matrix with no off-diagonal mass — `N_eff = N`: roster and effective count coincide, and that happens only at exact mutual orthogonality. At `PR = 1` — every vector the same — `N_eff = 1`, and the size of the roster has left the arithmetic entirely. **Everything the survey will ever report about `N` lives between those two poles, and the participation ratio is the coordinate that locates it.** When a sponsor asks why disagreement is averaged pairwise rather than some other way, cite the identity, not a plausibility argument.

Two distinctions you must keep separate, because collapsing them is the second most common survey error after the proxy trap:

Capacity `u` versus realised `u`. The probe set measures what the category set is *capable* of — how non-parallel the disciplines are when all of them speak. The period-by-period observable measures what actually *happened* — which categories were present, which spoke, which were overruled before speaking. These are different numbers and the gap between them is one of the cheapest things you will ever fix. The Worked Example below has capacity `U = 0.6972` and realised `U = 0.6019`, and the entire difference comes from sessions run with fewer than the full board present.

Present versus counted. A category that is in the room but silent contributes `x_i` to nothing. Under Encoder A you must define, in advance, what counts as participation — a recorded disposition, a spoken objection, a signed line. Attendance is not participation and must never be encoded as it.

24.4 Choice 3 — the order observable

This is the choice that decides whether your survey is worth anything. The order observable `Y_t` is a countable, durable quantity that stands in for `Q_t`.

A valid Y has five properties. Test each one explicitly and record the answer.

Property	Test	Why it matters
Durable	Does it persist for many periods with no input?	A stock, not a flow
Countable	Can two auditors independently produce numbers within a declared tolerance?	Otherwise `delta` and `eta` are noise
Decaying	Does it measurably decline under zero input?	If it never decays, `delta = 0` and the model has no fixed point
Independent of U	Different instrument, different record, different time?	Paper E §6.3 — see §24.15
Denominable	Can cost be expressed in the same units?	Otherwise h cannot be computed at all

Good order observables, by system type:

System	Candidate Y	Units
Review office / standards body	live precedent rules in the standing manual	rules
Research group	reproducible protocols in the group handbook	protocols
Software organisation	modules with a named owner and a passing test suite	modules
Library reading room	seat-hours of optional occupancy per week	seat-hours
Maintenance regime	assets with a current, executed maintenance record	assets
Co-housing block	standing agreements in force and observed	agreements
Manufacturing cell	validated work instructions in use on the floor	instructions

Bad order observables, and why:

- **Headcount.** Does not decay in the right way, is not produced by interaction, and is an input to N — it is on the wrong side of the equation.
- **Revenue.** A flow, not a stock, and driven by exogenous demand.
- **Engagement scores.** Not durable; they are measured, not held.

- **Meeting count.** This is U's support, not Q. Counting meetings as order is the purest form of the trap in §24.15.
- **Any index the organisation reports upward.** See §24.15 and Goodhart's Law.

24.5 Choice 4 — the time scale

One period is the unit of the recursion. Choosing it wrongly does not bias your estimates; it destroys your ability to detect anything.

Three constraints bound the choice.

1. **The period must be at least as long as the conversion lag.** If a decision made in a meeting takes six weeks to reach the manual, a one-week period puts the cause and the effect in different rows and the fit will find nothing.
2. **The period must be short enough that `delta` per period is small** — as a practical rule, `delta <= 0.3` per period, so that the linear approximation of a decaying stock is not being asked to cover a near-total loss in one step.
3. **The period must be long enough that the input term is visible above measurement noise.** This is the constraint nobody thinks about, and it has a precise structure. Audit error on a *level* does not accumulate: if you re-count the stock twice as far apart, the error at the endpoints stays the same size while the accumulated input doubles. **Coarsening the time scale raises the signal-to-noise of the input term roughly as the square root of the block length, and costs you degrees of freedom linearly.** There is an optimum, and you must choose it before you look at the answer (Chapter 25, §25.8).

Declare the period, the calendar alignment (which day the audit is taken), and the rule for handling periods with no interaction at all. Periods with `U_t = 0` are not missing data; they are the most informative periods you have, because they are where `delta` is identified.

24.6 Choice 5 — the cost observable, and the denomination rule

`C` must be denominated in the same units as `Y`. If you cannot denominate it, you cannot compute headroom, and the survey has failed. This is the single most frequently botched line in a real survey, because cost naturally arrives in money and hours while Y naturally arrives in rules, protocols, or modules.

There are exactly two honest routes.

Route 1 — direct destruction accounting. Count the Y-units destroyed or made unusable per period. In the Worked Example: rules struck by appeal reversal; rules orphaned because the only person who could apply them left; rules rendered unusable by contradiction with a newer rule. These are already in Y-units and require no conversion.

Route 2 — a declared conversion rate. For dissipation that is real but not directly destructive — coordination overhead, rework, interruption, emotional labour, compliance reporting — measure it in its natural unit and convert at a rate you declare in advance and justify. The honest way to set the rate is by *displacement*: measure how many Y-units the system produces per staff-hour of productive capacity, then charge dissipated hours at that rate. **A declared conversion rate is a badged modelling choice. An undeclared one is a fabrication.**

Cost categories to enumerate, from Chapter 5, with their usual Y-denominated form:

Category	Natural unit	Denomination route
Coordination overhead (Brooks, n(n−1)/2 paths)	hours	Route 2
Rework and re-decision	items	Route 1 if it strikes Y, else Route 2
Conflict and appeal	cases	Route 1
Turnover loss	departures	Route 1 (orphaned units)
Maintenance and compliance	hours	Route 2
Rent, commute, physical friction	money, minutes	Route 2
Interruption / context switching	switches	Route 2

Two rules that prevent the usual mistakes.

- **Do not double-count turnover.** Turnover appears in `delta` as decay of stored order *and* in `C` as a one-off destruction. Assign each departure to one or the other and write down which. The convention that works: routine attrition at the historical rate is in `delta`; departures above that rate, and any departure that orphans a specific Y-unit, are in `C`.
- **Do not omit the cost the organisation does not book.** The attention tax — the hours spent by people who are not in the meeting, reading about the meeting — is almost never in anyone's ledger, and it is often the largest single line. If you leave it out, your h will be optimistic and your intervention will fail for a reason you did not measure.

24.7 Estimating delta

`delta` is the fraction of stored order lost per period with no maintenance and no input.

Method D1 — the zero-input episode (best). Find a stretch of n periods in which `U_t = 0` and `C_t = 0`: a shutdown, a hiring freeze, a suspended committee, a closed season, a strike, a vacancy. Over that stretch the law reduces to `Q_{t+n} = (1 − delta)^n Q_t`, so

```
delta = 1 − (Y_end / Y_start)^(1/n)
```

This is the only clean estimator, and finding such an episode in the record is the single highest value hour of the whole survey. Institutions have more of them than they admit: budget freezes, pandemic closures, leave-of-absence gaps, the six months between the old director and the new one.

Method D2 — the cohort decay curve. Tag a cohort of Y-units created in one period and count how many are still live after 1, 2, 4, 8, 16 periods. Fit `ln(survivors)` against time; the slope is `ln(1 − delta)`. This is more work than D1 but works without a zero-input episode and it also tells you whether decay is exponential at all. If the survival curve is not straight in log space, the constant-`delta` assumption is wrong for your system and you must say so.

Method D3 — the maintenance ledger (fallback). If the organisation records maintenance effort per Y-unit, `delta` is approximately the fraction of units that would lapse in one period if that effort stopped. This is an estimate from testimony, not measurement, and must be labelled so.

Convert `delta` into the two numbers a non-technical reader can use:

```
half-life  t_half = ln 2 / ( -ln(1 - delta) )       periods
time constant  tau = 1 / delta                       periods
settling to 95% of a new equilibrium: about 3 tau    periods
```

delta	half-life (periods)	tau (periods)	3 tau (periods)
0.003	230.7	333.3	1000
0.005	138.3	200.0	600
0.01	69.0	100.0	300
0.02	34.3	50.0	150
0.05	13.5	20.0	60
0.10	6.6	10.0	30
0.20	3.1	5.0	15
0.30	1.9	3.3	10
0.50	1.0	2.0	6

A low `delta` is not unambiguously good. It buys durability and it charges latency: a system with `delta = 0.003` per week takes a thousand weeks to settle to a new equilibrium. Chapter 26 makes this the reason you promise increments and never promise equilibria.

24.8 Estimating eta

`eta` is Y-units of durable order created per nat of surplus input. It is the parameter most organisations have never considered and the one most often closest to zero.

Method E1 — the attribution audit (recommended). Choose a window of `w` periods in which the system was operating normally. For each Y-unit created in that window, ask a single question with a written answer: *which interaction produced this, and is that interaction in the record?* Count only the units with an affirmative, documented answer. Then

```
eta_hat = (attributable Y-units created in the window) / (sum of F(U_t) over the window)
```

The denominator is in nats; the numerator is in Y-units; `eta` is Y-units per nat.

The audit is deliberately strict. A unit created by one person alone at a desk is not attributable to interaction and does not count. A unit whose origin nobody can name does not count. **The strictness is the point: `eta` is supposed to measure the fraction of interaction that becomes structure, and a permissive audit measures the fraction of structure that had a meeting nearby.**

Method E2 — the capture-rate decomposition. For systems where the attribution audit is too expensive, decompose:

```
eta ~= (fraction of interactions that produce a written decision)
     x (fraction of written decisions that reach the durable medium)
     x (Y-units per decision that arrives)
```

Each factor is separately observable and each names its own fix. A meeting culture with excellent minutes and no publication path has a high first factor and a near-zero second one.

Method E3 — the natural experiment. If the organisation changed its capture practice at a known date — introduced a decision log, hired a technical writer, mandated a template — estimate

`eta` separately before and after. This is the cheapest way to learn what an `eta` intervention is worth in your own system, and Chapter 26 will ask you for exactly that number.

24.9 Estimating N

`N` is the number of genuinely distinct categories the system spans. It appears in the ceiling `ln N` and therefore in every design inequality.

Procedure:

1. **Enumerate candidate categories.** Disciplines, training lineages, functions, sites, tenures, client types — whatever the system itself treats as a kind.
2. **Merge duplicates.** Two candidates merge if their encoded vectors have `u_ij < 0.10` — that is, if they disagree on less than about 2.6% of probe cases under Encoder A. Two categories that never disagree are one category with two names, and counting them twice inflates `ln N` while contributing nothing.
3. **Drop non-participants.** A category that has no route into the interaction is not in `N`. A discipline that is consulted after the decision is not in `N`.
4. **Count what survives.** That is `N`. Record the merge and drop decisions with reasons.

The marginal value of `N` is brutal and should be shown to anyone who proposes hiring their way out of a headroom problem:

N	ln N	gain over N-1
2	0.6931	0.6931
3	1.0986	0.4055
4	1.3863	0.2877
5	1.6094	0.2231
6	1.7918	0.1823
8	2.0794	0.1335
10	2.3026	0.1054
12	2.4849	0.0870
16	2.7726	0.0645
20	2.9957	0.0513
30	3.4012	0.0339
50	3.9120	0.0202
100	4.6052	0.0101

Against Brooks' communication paths `n(n−1)/2`:

n	paths	paths added by the n-th member
3	3	2
4	6	3
5	10	4
6	15	5
8	28	7
10	45	9

n	paths	paths added by the n-th member
12	66	11
15	105	14
20	190	19
26	325	25

Surplus grows logarithmically and coordination cost grows quadratically. That is the whole team sizing problem in one line, and the survey's job is to find where your system sits on it.

24.10 Estimating u

With the encoder fixed (§24.2) and the aggregator declared (§24.3), u is arithmetic. What takes judgement is the instrument that feeds it.

Instrument U1 — the probe set (capacity). 8 to 16 real cases, chosen to span the system's actual workload, each with a forced binary disposition. Administered independently and in writing *before* any discussion, because discussion is exactly the mechanism by which cascades collapse u (Bikhchandani, Hirshleifer and Welch 1992). Administer it annually; u is not a property you acquire, it is a quantity that decays.

Instrument U2 — the disagreement audit (realised). From the decision record, count for each period the pairs of categories that registered opposed positions before the decision closed. This is a weaker signal than the probe set — a category can be silent because it agrees or because it has learned not to speak — but it is available for every period retrospectively, which the probe set is not.

Instrument U3 — error decorrelation (strongest, where available). Where the system makes resolvable predictions, compute the correlation of signed errors between categories. Decorrelated error is the operational definition of independent judgement. It is also the only instrument that cannot be gamed by disagreeing performatively.

Instrument U4 — the collision test (fast, from Book One). Ask each category, independently and in writing, for the three things that would most change their mind. Count the distinct answers. A category set that returns the same three answers has u near zero, whatever its demographics.

The effective-count table for $N = 5$, which is the arithmetic that makes u legible to a manager:

u	$N_{eff} = 1 + 4u$	$F(u) = \ln N_{eff}$
0.00	1.000	0.0000
0.10	1.400	0.3365
0.20	1.800	0.5878
0.30	2.200	0.7885
0.40	2.600	0.9555
0.50	3.000	1.0986
0.5715	3.286	1.1897
0.60	3.400	1.2238
0.6972	3.789	1.3320
0.80	4.200	1.4351

u	N_eff = 1 + 4u	F(u) = ln N_eff
0.90	4.600	1.5261
1.00	5.000	1.6094

Five reviewers at `u = 0` are one reviewer. The survey's job is to report the effective count, not the roster.

24.11 Estimating C

Follow §24.6's denomination rule and build the tally as a line-item table, one line per category, each with its measurement method and its conversion (if Route 2). The tally is the only part of the survey where the organisation will actively argue with you, because every line is somebody's work.

Discipline that keeps the tally honest:

1. **Every line has a source document.** Appeal register, exit interviews, ticket system, time records, the minutes.
2. **Every Route 2 line shows its conversion rate and the justification for it, in the tally.**
3. **The tally is reconciled against the series.** If your line-item `C_hat` differs from the implied cost in the observed trajectory by more than your declared tolerance, one of them is wrong. In the Worked Example the line-item tally came to 0.63 `Y`-units per week against a series-implied 0.6275 — agreement to 0.0025, which is well inside tolerance and is the only reason the tally may be used.
4. **Slack is not cost.** Queueing theory is unambiguous that utilisation near 1 makes waiting time explode; unused capacity is a cost-control device, not a dissipation. Booking slack as `C` inverts the design conclusion.

24.12 Estimating Q

`Q` is simply `Y` at the survey date, once you have satisfied yourself that the count is the count.

Two corrections are almost always needed, and they are different in kind.

Correction 1 — the liveness audit. Units that are formally in the register but which nothing and nobody has used in the last `k` periods are not stored order; they are stored paper. Define `k` in advance, audit, and remove them from `Y`. **Striking units that were already inert is not shedding `Q`; it is correcting the measurement.** The distinction matters enormously in Chapter 26, because correcting the measurement changes the number and not the world, whereas shedding `Q` is an intervention with real effects.

Correction 2 — the applier audit. A unit that survives but which no remaining person can apply is orphaned. Orphaning is a real loss and belongs either in `delta` (if routine) or in `C` (if caused by a specific departure) — but the unit itself must come out of `Y` either way.

Record the pre-correction count, both corrections, and the post-correction `Q`, so a second engineer can see what you removed and why.

24.13 Instruments

The five choices tell you what to measure. These are the instruments that take the measurements.

Direct observation, Whyte's method. William H. Whyte's *The Social Life of Small Urban Spaces* (1980) established the method this manual takes as canonical for physical systems: go, sit, watch, count, and record what people actually do rather than what they say they do. For an organisation the equivalent is attending the interaction without a role in it and tallying: who spoke, in what order, for how long, who was interrupted, what was written down, and by whom. **Whyte's finding that people sit where there is something to sit on generalises: systems do what their fixtures make easy.** An hour of direct observation routinely overturns a month of interviews.

Time-lapse and the record. Whyte used time-lapse film because occupancy patterns are invisible at human tempo. The organisational equivalents are the commit log, the calendar export, the ticket system, the door-badge record, the room-booking data. These are unglamorous and they are the only sources in most organisations that were not written to be read by management.

The decision log. The single most useful artefact to install if it does not exist, and the single most useful to read if it does. Each entry: date, question, who was present, positions recorded before discussion, decision, and the durable artefact produced. A decision log gives you `U_t` and `eta` at once, and a system that has kept one for two years can be surveyed in a week.

The disagreement audit. A structured pass over the decision record, coding each decision for whether any category recorded an opposed position, and whether that position survived to the written outcome. The ratio of the second to the first is one of the best single diagnostics in the manual: an organisation where dissent is recorded and never survives has a governance problem that u alone will not reveal.

Turnover and tenure records. These feed `delta` directly and `C` through orphaning. Get the distribution, not the average; a department with a stable core and a churning periphery has a very different `delta` profile from one with uniform three-year tenure.

The desire-line survey — the cheap instrument that measures the same gradient. Chapter 42 derives an identity that changes the economics of this chapter: **`grad h = −grad C`.** Since `h = eta ln N − delta Q − C` and only `C` varies with the local design choice under consideration, the direction of steepest increase in headroom is exactly the direction of steepest decrease in cost — and cost, at the point where it is felt, is friction. Book One's Node 02 called the same object the desire line: the path worn across the grass because the paved route costs more than it is worth.

The operational consequence is that a cheap habit survey IS a headroom measurement, in the direction if not in the level. Ask twenty people, in writing, where the system makes them work around it — which meeting they skip, which form they fill in wrong on purpose, which route they take instead of the one drawn. Tally the answers. That tally points along `−grad C`, which is `grad h`, and it costs an afternoon rather than the six-week instrument programme the rest of this chapter specifies.

What it does not give you is the sign or the magnitude of h itself. **A gradient tells you which way is uphill; it does not tell you that you have the budget to climb.** So: run the desire-line survey first, always, to locate the lever; run the full survey when you need to know whether you may pull it. Chapter 26 §26.16 makes that distinction a gate.

Star and Ruhleder's warning. Infrastructure becomes visible only on breakdown (1996). This

has a specific consequence for surveying: **you will systematically under-measure everything that is working.** The maintenance regime that never fails will have no tickets, no incidents, no records and no budget line. Ask explicitly, for each Y-unit class, what keeps it alive, and expect the answer to be a person rather than a process.

24.14 Computing headroom and reading the verdict

With the six numbers in hand:

```
eta ln N        the best case: what conversion could deliver at the ceiling
delta Q         the maintenance draw: what the existing stock costs to hold
C_hat           the dissipation draw
h = eta ln N − delta Q − C_hat
```

Read it as a budget. `eta ln N` is the income at maximum geometric efficiency. `delta Q` is the mortgage on what you already own. `C_hat` is the operating expense. `h` is what is left.

Then compute the *realised* one-step change, which is a different and equally important number:

```
one−step = eta F(U_realised) − delta Q − C_hat
```

`h` uses the ceiling `ln N`; the one-step uses what actually happened. The gap between them is the **realisation deficit** — headroom that exists and is not being collected. Chapter 26 treats these as two different classes of lever and it matters that you report both.

One correction on `F(U_realised)`, and it is a correction most surveys get wrong. `U` is not constant across the periods in your series; it is a distribution. What the accumulator integrates is `sum of eta F(U_t)`, so the operative quantity is the **mean of F**, not `F` of the mean. And because `F` is strictly concave (Paper A, derived result 3.3), Jensen's inequality makes those two numbers different in a known direction:

```
mean of F(U_t)  <=  F( mean of U_t )       with equality only if U is constant
```

Compute `F` period by period and average the results. Never average `U` and then take `F` once. The gap is the concavity penalty and it is not small: in the Chapter 26 worked office it is 0.0761 nats against an `F(mean U)` of 1.2262 — **6.2% of the surplus, lost purely to the variability of the geometry.** A survey that reports `F(mean U)` overstates the system's realised surplus by that much, and every downstream number — the one-step, `Q_inf`, the feasibility check of Chapter 26 §26.7 — inherits the error. Report both figures on the output page, label the difference as the concavity penalty, and test against the mean of `F`.

Finally compute the steady state, `Q_inf = (eta F(U) − C)/delta` (Theorem 3.5). **If `Q_inf` is negative, the theorem does not hide the failure mode: the regime is unsustainable, or the linear model must be replaced.** A negative `Q_inf` in a survey output is not a rounding artefact and must never be reported as "approximately zero". It is the model telling you, in its own terms, that the system as configured is being driven to the boundary of its own validity.

The cheapest decisive prediction the survey can make, and it should be made every time. `u` decays; Chapter 11 gives the mechanisms and §24.10 says plainly that `u` is not a property you acquire. Hold `D`, `p_bar`, `N` and the budget constant, let `u` decay geometrically at rate `lambda` per period, and the doctrine-to-data ratio `r` of Chapter 48 crosses 1 at

```
t* = ln(u*/u_0) / ln(1 − lambda)          periods
```

t∗ has no free parameters. Every quantity in it — the current u_0, the crossing level u∗, the decay rate lambda — is already on the survey form, and nothing has to be fitted, tuned or assumed. In the worked case of Chapter 48 it comes to about 8.4 quarters. That makes it **the most decisive and least expensive falsification available anywhere in this part of the manual**: write the date down, wait, and look. A system that sails past t∗ with r still below 1 has falsified either the decay estimate or the model, and either finding is worth more than another season of fitting.

And a hard rule about combining surveys, because the arithmetic is unforgiving. Do not aggregate r across sites, divisions or cohorts. Distortion does not average. Two sites distorted in opposite directions — one over-framed, one under-framed — produce an aggregate D/P in which the two errors cancel, and the cancellation is not a reconciliation; it is a deletion. In the two-site example of Chapter 48 the aggregate **overstated bankable input by 4.26x** while each site taken alone was correctly diagnosed. **Survey each site to its own r and report the vector, never the mean.** An organisation with one site at r = 2 and one at r = 0.5 does not have an organisation at r = 1; it has two different diseases and no healthy site at all.

24.15 The order-proxy trap — the most common measurement error

Everything else in this chapter is craft. This section is the part that decides whether the craft was worth anything.

Paper E §6.3: you must justify that your order proxy Y is not merely a repackaging of the input U. If it is, the estimation model of Chapter 25 will appear to fit — often superbly — and will have measured nothing about the world.

Why it happens is not stupidity. It happens because the most available records are the interaction records. The minutes exist. The attendance list exists. The ticket system exists. So the engineer in a hurry defines Y as "items dispositioned per session" or "decisions logged per month" or "a weighted index of board activity" — and every one of those is a function of the interaction, not a store left behind by it. The equation then reads, in effect, U = a U + b U, and the fit recovers an identity.

There are two distinct species and they need different tests.

Species 1 — contemporaneous repackaging. Y_t is computed from the same records as U_t, at the same time. Example: Y = items dispositioned this session; U = disagreement observed this session. Both are functions of how busy the board was.

Species 2 — filtered repackaging. Y_t is a smoothed or accumulated version of U itself. Example: Y = a rolling index of the board's disagreement score, published weekly by the clerk who writes the minutes. This species is far more dangerous because a filtered input is *genuinely* predicted by the input, so the model fits almost exactly and — as Chapter 25 demonstrates with real numbers — **passes every statistical baseline including the shuffled-input control.**

Run all four tests. Record the answers in the survey. Two of them are structural and must be run before you take a single measurement.

Test 1 — the algebraic derivation test (structural). Write out, as a formula, how Y_t is computed from source fields. Write out how U_t is computed. Compare the source fields. If any field appears in both, or if Y's formula contains U's formula, you have repackaging. Declared tolerance: **zero shared fields.** This test costs ten minutes and catches most cases.

Test 2 — the instrument-independence test (structural). Y must be recorded by a different person, into a different medium, on a different schedule, from U. In the Worked Example, U comes from the board's minutes and Y comes from the code administrator's register; the code administrator does not attend the board. If the same clerk produces both, the test fails regardless of how the formulas look, because the two numbers will share the clerk's judgement.

Test 3 — the zero-input persistence test (empirical, decisive). Find a period with `U = 0` and look at Y. **A stock persists and decays slowly; a repackaged input collapses to zero and springs back.** This is the strongest empirical test in the manual and it is available whenever you have a shutdown episode — the same episode you needed for `delta`. In the Worked Example, the true proxy went 206 to 186 across a 34-week freeze and continued from 186 when the board resumed. The repackaged proxy fell from 82.1 to 0.1 and returned to 31.6 in a single week.

Test 4 — the ceiling test (arithmetic, decisive). The law itself bounds how fast a genuine stock can grow. From Corollary 4.2, the largest possible one-step increment is `eta ln N − delta Q − C <= eta ln N`. In the Worked Example that ceiling is `0.5015 x 1.6094 = 0.8071` Y-units per week. The repackaged proxy's single-week jump of 31.1 units is **38.5 times the theoretical maximum**. A quantity that can move faster than the conversion capacity of the system that allegedly produced it is not stored order. **The law is not only a model; it is a consistency check on your own instrument.**

Two further tells worth recording, which do not prove repackaging but reliably accompany it:

- **A fitted a implying an implausible half-life.** In the Worked Example the fitted a for the filtered repackaging was 0.7038, implying a half-life of 1.97 weeks, against a directly measured half-life of 224 weeks — a discrepancy of 114-fold. **An institutional stock with a two-week half-life is not an institutional stock.**
- **A fitted `b` far outside the audited `eta`.** The same fit returned `b = 20.2` against an attribution-audited `eta` of 0.5015 — a factor of 40. `b` is supposed to be `eta`. When it is not in the same order of magnitude, the two are not measuring the same thing.

The shuffle will not save you, and here are the three measured reasons. The four-baseline suite of Chapter 25 ends with the shuffled-input control, and engineers reach for it as the general-purpose guarantee that a fit means something. It is not one. Chapters 55 and 81 measured its behaviour by simulation rather than asserting it, and it fails in three distinct ways that the surveyor must know before designing the instrument, because two of the three can only be fixed at design time.

(1) It is anti-conservative as ordinarily run. Against a nominal alpha of 0.05 across 500 null trials the measured false-rejection rates were **0.108 for iid permutation (2.2x nominal), 0.120 for circular shift (2.4x nominal), and 0.060 for block permutation** (near nominal). A surrogate test taken off the shelf roughly doubles your false-positive rate. **The remedy is not to abandon it but to calibrate it by simulation on your own study's design** — your series length, your autocorrelation, your block structure — before you use it. Chapter 55 makes that calibration a requirement.

(2) It cannot see repackaging. The filtered repackaging constructed in Chapter 25 passed all four baselines, the shuffle included, at `p = 0.0005`. It was caught by Test 4 above and by nothing else: its single-step jump was **38.5 times the theoretical maximum** set by Corollary 4.2. **Repackaging is a structural fault and only a structural test detects it.**

(3) It cannot see confounding by a shared slow driver. Where U and Y are both driven by a third

slow variable — a funding cycle, a hiring wave, a market season — a 4,000-trial simulation found the shuffled-input control **passing 98.35% of the time**. The same simulation reports the control working honestly in the clean case, at 48.17% detection, which is what makes the 98.35% figure a measurement rather than a complaint.

The consolidated rule: more surrogates do not fix any of this. Design controls do. Blinding, channel separation, prospective sampling and a yoked control are the instruments that answer (2) and (3), and every one of them is a decision you make in §24.2 through §24.6, before a single number exists. **This is why the repackaging tests sit in the survey chapter and not in the statistics chapter: by the time you are permuting columns, the instrument is already built.**

The Specification

Numbered, checkable requirements. A survey that does not satisfy all of these is not a survey and must not be used as the basis for an intervention under Chapter 26.

S24.1 The system boundary shall be stated in writing, including at least one named thing that is deliberately excluded and the reason for excluding it.

S24.2 All five operational choices of Paper E §5.1 shall be recorded before any measurement is taken, and shall be dated.

S24.3 The state encoder shall be specified to the level of "a second engineer could produce the same vectors": the probe set, the response format, the normalisation, and the administration protocol.

S24.4 The probe set shall contain at least 8 real cases drawn from the system's actual workload and shall be administered independently and in writing before discussion.

S24.5 The interaction aggregator shall be declared and badged as a modelling choice. If the mean pairwise aggregator is used, its identity with the participation-ratio effective count shall be noted so the choice is traceable to a standard diversity measure.

S24.6 Capacity u and realised u shall be reported as two separate numbers with their two separate instruments. Reporting only one is a specification violation.

S24.7 The order observable shall be tested against all five properties of §24.4, with the answers recorded.

S24.8 The four repackaging tests of §24.15 shall be run and their results recorded. Tests 1 and 2 shall be run before measurement begins.

S24.9 Test 1 shall be passed with zero shared source fields between `Y` and `U`.

S24.10 Test 2 shall be passed with a named different recorder, a named different medium, and a named different schedule.

S24.11 Where a zero-input episode exists in the record, Test 3 shall be run and its result plotted.

S24.12 Test 4 shall be run for the largest single-period increase in `Y` observed in the series. Any increase exceeding `eta ln N` shall be investigated before the survey is released.

S24.13 The time scale shall be declared with its calendar alignment and shall satisfy: period >= conversion lag; `delta` per period <= 0.3; and a stated rationale for the signal-to-noise trade-off of §24.5.

S24.14 Periods with `U_t = 0` shall be retained in the series and flagged, never dropped.

S24.15 `C` shall be denominated in Y-units. Every Route 2 line shall show its conversion rate and the justification for it.

S24.16 The attention tax shall appear as an explicit line in the cost tally, even if the estimate is coarse. A tally with no attention-tax line is incomplete.

S24.17 Turnover shall be assigned to `delta` or to `C` by a written rule, and shall not appear in both.

S24.18 Slack shall not be booked as cost.

S24.19 `delta` shall be estimated by Method D1 where a zero-input episode exists. Where D2 is used, the log-survival plot shall be inspected for linearity and non-linearity reported.

S24.20 `delta` shall be reported with its half-life and time constant in periods and in calendar time.

S24.21 `eta` shall be estimated by an attribution audit in which each counted Y-unit has a documented originating interaction. Unattributable units shall be excluded and their count reported.

S24.22 `N` shall be produced by explicit enumeration, merge, and drop steps, with each merge and drop recorded with its reason.

S24.23 Categories with pairwise `u < 0.10` shall be merged.

S24.24 `Q` shall be reported pre-correction, with the liveness correction and the applier correction shown separately, and post-correction.

S24.25 The line-item `C_hat` shall be reconciled against the cost implied by the observed series, and the residual reported. A residual exceeding the declared tolerance blocks release.

S24.26 `h`, the realised one-step change, and `Q_inf` shall all be computed and reported.

S24.27 A negative `Q_inf` shall be reported as negative, with the Theorem 3.5 failure clause quoted.

S24.28 Every number on the survey output page shall carry a provenance note naming the instrument and the date range.

S24.29 The survey shall state, in one sentence each, the three things that would most change its conclusions if measured differently.

S24.30 The survey shall not recommend an intervention. Recommendation is Chapter 26's function and requires the preregistration of Chapter 25 first.

The Survey Worksheet

Printable. Seven forms. Fill in order; do not skip forward.

```
============================================================
FORM A — SCOPE AND THE FIVE CHOICES                  Date: ____
============================================================
A1  System name ..........................................
A2  Boundary: what is IN .................................
```

```
A3  Boundary: what is OUT, and why .......................
A4  Surveyor(s) ..........................................

CHOICE 1 — STATE ENCODER
A5  Encoder type  [ ] A disposition  [ ] B error  [ ] C practice
A6  Probe set: m = ____ cases. Listed at .................
A7  Response format ......................................
A8  Administration: independent? [ ] Y [ ] N   written? [ ] Y [ ] N
A9  Normalisation ........................................

CHOICE 2 — INTERACTION OBSERVABLE
A10 Aggregator ...........................................
A11 BADGE: aggregator is a declared modelling choice  [ ] initialled
A12 Participation rule (what counts as taking part) ........
A13 Capacity instrument ..................................
A14 Realised instrument ..................................

CHOICE 3 — ORDER OBSERVABLE
A15 Y = ..................................................
A16 Units ................................................
A17 Register / medium ....................................
A18 Recorder (name and role) ..............................
A19 Audit schedule .......................................

CHOICE 4 — TIME SCALE
A20 One period = .........................................
A21 Calendar alignment ....................................
A22 Conversion lag measured at ____ periods
A23 Rationale for period length ...........................

CHOICE 5 — COST OBSERVABLE
A24 Route 1 categories ....................................
A25 Route 2 categories ....................................
A26 Declared conversion rate ____ Y-units per ____
A27 Justification for the rate ............................

================================================================
FORM B — THE REPACKAGING AUDIT  (Paper E section 6.3)
================================================================
B1  Y source fields:  ....................................
B2  U source fields:  ....................................
B3  TEST 1 shared fields count = ____      PASS [ ]  FAIL [ ]
B4  Y recorder _________  U recorder _________ different? [ ]
B5  Y medium _________    U medium _________   different? [ ]
B6  Y schedule _________  U schedule _________ different? [ ]
B7  TEST 2                                   PASS [ ]  FAIL [ ]
B8  Zero-input episode found? [ ] Y [ ] N   periods ____ to ____
B9  Y at episode start ____  at end ____  after resumption ____
```

```
B10 TEST 3: does Y persist and decay smoothly? PASS [ ] FAIL [ ]
B11 Largest single-period rise in Y = ____
B12 Ceiling eta*lnN = ____   ratio = ____
B13 TEST 4                                PASS [ ]  FAIL [ ]
B14 IF ANY TEST FAILS: STOP. Redefine Y. Do not proceed.

================================================================
FORM C — CATEGORIES AND GEOMETRY
================================================================
C1  Candidate categories enumerated:  ......................
C2  Merges (pair, u_ij, reason):  ..........................
C3  Drops (category, reason):  ............................
C4  N = ____         ln N = ____
C5  PAIRWISE TABLE
    pair | disagree k/m | d | cos = 1-2d | u = 1-cos^2
    -----+--------------+---+------------+------------
         |              |   |            |
C6  U capacity = mean pairwise u = ____
C7  N_eff capacity = 1 + (N-1)U = ____   F = ln N_eff = ____
C8  U realised (from period record) = ____
C9  F realised = ____
C10 Realisation deficit F_cap - F_real = ____

================================================================
FORM D — DELTA
================================================================
D1  Method [ ] D1 zero-input [ ] D2 cohort [ ] D3 ledger
D2  Episode / cohort identified ............................
D3  n periods = ____   Y_start = ____   Y_end = ____
D4  ratio = ____   delta = 1 - ratio^(1/n) = ____
D5  half-life = ln2 / (-ln(1-delta)) = ____ periods = ____ calendar
D6  tau = 1/delta = ____ periods    3*tau = ____ periods
D7  If D2: is log-survival linear? [ ] Y [ ] N   note ........

================================================================
FORM E — ETA
================================================================
E1  Window: periods ____ to ____   (w = ____)
E2  Y-units created in window = ____
E3  Of these, ATTRIBUTABLE with documented origin = ____
E4  Unattributable (excluded) = ____
E5  sum of F(U_t) over window = ____ nats
E6  eta = E3 / E5 = ____ Y-units per nat
E7  Capture decomposition (optional):
    interactions producing a written decision ....... ____
    written decisions reaching the durable medium ... ____
    Y-units per arriving decision .................. ____
```

```
================================================================
FORM F — COST
================================================================
line | category | natural measure | route | rate | Y-units/period
-----+----------+-----------------+-------+------+---------------
  1  |          |                 |       |      |
  2  |          |                 |       |      |
  3  |          |                 |       |      |
  4  |          |                 |       |      |
  5  |          |                 |       |      |
F1  Attention-tax line present? [ ] Y  (required)
F2  C_hat TOTAL = ____ Y-units/period
F3  Series-implied cost = ____   residual = ____
F4  Residual within declared tolerance ____ ? [ ] Y [ ] N
F5  Turnover assigned to [ ] delta [ ] C   rule: ............
F6  Slack booked as cost? [ ] N  (must be N)

================================================================
FORM G — Q, HEADROOM, VERDICT
================================================================
G1  Y register count at survey date = ____
G2  Liveness correction (inert units removed) = -____
G3  Applier correction (orphaned units removed) = -____
G4  Q = ____
G5  eta * ln N = ____
G6  delta * Q = ____
G7  C_hat = ____
G8  h = G5 - G6 - G7 = ____
G9  VERDICT  [ ] h < 0 unsustainable   [ ] h = 0 hold
             [ ] h > 0 growth available, r_max = ____
G10 realised one-step = eta*F_real - delta*Q - C_hat = ____
G11 realisation deficit = G8 - G10 = ____
G12 Q_inf = (eta*F_real - C_hat)/delta = ____
G13 Q_inf negative? [ ] Y -> quote Theorem 3.5 failure clause
G14 Three measurements that would most change this conclusion:
    1 ..........................................................
    2 ..........................................................
    3 ..........................................................
G15 Surveyor signature _______________  Date __________
```

The Failure Modes

F24.1 — THE ORDER PROXY THAT WAS REALLY THE INPUT. *Description.* Y is a function of U, so the model fits an identity. *Damages.* Everything. The survey has no empirical content. *Signature.* Fitted b far outside the audited `eta`; fitted a implying an implausible half-life; Y collapses to zero during a shutdown; single-period jumps exceeding `eta ln N`. *Detection.* The four tests of

§24.15. Tests 1 and 2 before measurement; 3 and 4 on the series. *Fix.* Redefine Y into a different medium with a different recorder. There is no statistical remedy; this is a definitional error and only a definitional change repairs it.

F24.2 — THE MISSING DENOMINATOR. *Description.* `C` is measured in hours or money while Y is measured in rules or modules, and h is computed anyway by implicitly treating them as commensurate. *Damages.* `C`, and therefore h. *Signature.* The headroom calculation has mixed units. The number is arithmetically well-formed and physically meaningless. *Detection.* Write the units next to every term in `h = eta ln N − delta Q − C`. They must match. *Fix.* §24.6. Route 1 where possible; a declared and justified Route 2 rate otherwise.

F24.3 — THE SNAPSHOT SURVEY. *Description.* `Q` is measured once. No series, therefore no `delta`, therefore no h. *Damages.* `delta`, and by omission the verdict. *Signature.* The survey reports a level and a judgement about whether it is "healthy". *Detection.* Ask for the second measurement date. If there is not one, this is the fault. *Fix.* Reconstruct a retrospective series from the register's own history. Most durable registers carry creation and strike dates and can be replayed.

F24.4 — THE MISSING ZERO-INPUT EPISODE. *Description.* No shutdown in the record, so `delta` is estimated by testimony. *Damages.* `delta`, and through `delta Q` the whole budget. *Signature.* `delta` reported as a round number with no derivation. *Detection.* Check Form D for a named episode. *Fix.* Method D2, the cohort decay curve, which needs no episode. Failing that, run a deliberate, bounded, reversible zero-input trial on a subset — Chapter 26's pilot rules apply.

F24.5 — ATTENDANCE ENCODED AS PARTICIPATION. *Description.* Categories present but silent are counted in `U_t`. *Damages.* u, upward. Inflates `F(U)` and hides the realisation deficit. *Signature.* Realised u equals capacity u exactly, which essentially never happens. *Detection.* Compare the participation rule in A12 against the actual record of who spoke. *Fix.* Require a recorded disposition. A category with no recorded position is not in that period's geometry.

F24.6 — THE PERMISSIVE ATTRIBUTION AUDIT. *Description.* `eta` is computed by dividing all new Y-units by all surplus, without checking which units actually came from interaction. *Damages.* `eta`, upward, often by a large factor. *Signature.* `eta` implausibly high relative to the observed trajectory; the fitted b of Chapter 25 comes out far below the audited `eta`. *Detection.* Recount with the strict rule: documented originating interaction or excluded. *Fix.* Redo Form E. Report the unattributable count; it is itself a useful diagnostic, because a large unattributable fraction means the system's order is being made somewhere you are not looking.

F24.7 — CATEGORY INFLATION. *Description.* `N` counted from the org chart rather than from the geometry, so relabelled duplicates are counted separately. *Damages.* `N` upward, and therefore `eta ln N` and h upward. *Signature.* Two categories in the pairwise table with `u < 0.10` and separate rows. *Detection.* Form C, the merge step. *Fix.* Merge. Then note that the merge has *reduced* your measured headroom, which is the survey doing its job.

F24.8 — THE UNRECONCILED TALLY. *Description.* The cost line items are summed and never checked against the series. *Damages.* `C`, in either direction. *Signature.* Form F3 blank. *Detection.* Compute the cost implied by the observed trajectory given `delta`, `eta` and `F(U)`, and compare. *Fix.* Reconcile. A large residual usually means an omitted cost category, not an arithmetic error.

F24.9 — GOODHART CAPTURE OF THE PROXY. *Description.* Once Y is known to be measured, the organisation optimises Y directly. *Damages.* The validity of Y as a proxy for `Q`, progressively, over time. *Signature.* Y rises while every qualitative indicator falls. Newly created units are small,

uniform, and never applied. *Detection.* The liveness audit (§24.12) is the counter-instrument: Goodharted units are created and never used. *Fix.* Chapter 25 §25.10. In the survey, the immediate mitigations are to keep Y in a register that is not reported upward, to include the liveness correction as standing practice, and to rotate the proxy definition on a schedule known to the surveyor and not to the surveyed.

F24.10 — THE INVISIBLE MAINTENANCE REGIME. *Description.* Everything that works is invisible, so the survey under-measures both Q and the effort holding it (Star and Ruhleder 1996). *Damages.* Q downward and C downward — which can mask an unsustainable h behind two compensating errors. *Signature.* No records exist for the parts of the system nobody complains about. *Detection.* For each Y-unit class, ask what keeps it alive and who. Expect a person. *Fix.* Count the named people as a maintenance line in C and note the concentration risk. A maintenance regime resting on one un-deputised person is an orphaning event waiting for a date.

Worked Example

The system

Meridian County Development Review Office. A municipal development review office: 26 staff, five review disciplines, a weekly cross-discipline review board, and a standing review manual of written precedent rules that reviewers are required to apply.

Badge — this is a constructed composite, not a field study. The organisation is invented. The arithmetic is real, computed, and reproducible: the audit series is a simulation with declared generating parameters, produced so that every number in Chapters 24, 25 and 26 can be checked by the reader. **It must not be cited as evidence about any real review office.** Its function is to show the procedure operating end to end on numbers that behave the way real numbers behave — including the places where they refuse to cooperate.

The five choices

#	Choice	As declared for MCDRO
1	State encoder	Encoder A. Probe set of m = 12 real development applications from the previous two years, each with a forced approve/refuse disposition, recorded independently and in writing by each discipline lead before any discussion. Vector = the 12 signs, normalised.

#	Choice	As declared for MCDRO
2	Interaction observable	Mean pairwise `1 - (x_i . x_j)^2` over the disciplines that recorded a disposition in that week's board. Participation = a recorded disposition, not attendance. Badged as a declared aggregator.
3	Order observable	Y = count of **live precedent rules in the standing review manual**. A rule is live if it is in the register, not superseded, and has been applied at least once in the previous 156 weeks. Maintained by the code administrator, who does not attend the board.
4	Time scale	One week. Audit taken every Friday. Conversion lag measured at up to 6 weeks (drafting plus the 10-working-day publication requirement).
5	Cost observable	Route 1 for appeal strikes, orphaning, and contradiction withdrawals. Route 2 for rework and the attention tax, at a declared displacement rate.

Form B — the repackaging audit, run first

Test	Result
1 Algebraic derivation	Y fields: rule id, adoption date, supersession date, last-applied date. U fields: board date, discipline, recorded disposition. **Shared fields: 0. PASS.**
2 Instrument independence	Y recorded by the code administrator into the rule register, audited Fridays. U recorded by the board clerk into the minutes, at the session. Different person, medium, schedule. **PASS.**

Test	Result
3 Zero-input persistence	Weeks 121-154, board suspended by a budget freeze. Y: 206 at week 121, 186 at week 154, 186 at week 155 when the board resumed. Smooth monotone decay through the episode, continuous across resumption. **PASS.**
4 Ceiling	Largest single-week rise observed: 2 rules. Ceiling `eta ln N = 0.5015 x 1.6094 = 0.8071` per week; over the 6-week conversion lag the admissible burst is well above 2. **PASS.**

Only now does measurement begin.

Form C — categories and geometry

Five disciplines survive enumeration: planning (PLAN), structural and building (STRU), fire and life safety (FIRE), environmental (ENVR), transportation (TRAN). No merges: the smallest pairwise u is 0.3056, comfortably above the 0.10 merge threshold. No drops: all five have a route into the board.

`N = 5      ln N = 1.6094`

pair	opposed cases k/12	d	cos = 1 - 2d	u = 1 - cos^2
PLAN-STRU	3/12	0.2500	0.5000	0.7500
PLAN-FIRE	2/12	0.1667	0.6667	0.5556
PLAN-ENVR	5/12	0.4167	0.1667	0.9722
PLAN-TRAN	4/12	0.3333	0.3333	0.8889
STRU-FIRE	1/12	0.0833	0.8333	0.3056
STRU-ENVR	3/12	0.2500	0.5000	0.7500
STRU-TRAN	2/12	0.1667	0.6667	0.5556
FIRE-ENVR	3/12	0.2500	0.5000	0.7500
FIRE-TRAN	2/12	0.1667	0.6667	0.5556
ENVR-TRAN	4/12	0.3333	0.3333	0.8889

```
U capacity = 0.6972
N_eff capacity = 1 + 4(0.6972) = 3.7889
F capacity = ln 3.7889 = 1.3321 nats
```

Realised, from the 278 operating weeks of the audit series:

```
U realised = 0.6019      F realised = 1.1498 nats
realisation deficit = 1.3321 − 1.1498 = 0.1823 nats
```

The cause is visible in the record: **32.7% of board sessions ran with fewer than the full set of disciplines recording a disposition.** Those sessions have a mean u of 0.2071 against 0.7940 for full sessions. The office has five genuinely non-parallel disciplines and uses about four-fifths of them.

Note the FIRE-STRU pair at u = 0.3056. That is the closest thing to redundancy in the set and is exactly where a future reorganisation would be tempted to merge two roles. Merging them would drop N from 5 to 4 and cost ln 5 − ln 4 = 0.2231 nats of ceiling — which, at the office's current eta, is 0.112 rules per week, or about 5.8 rules a year, given away permanently.

Form D — delta

Method D1. The budget freeze of weeks 121-154 suspended the board entirely; no dispositions, no new rules, and — because intake was also halted — no appeals, rework, or contradiction maintenance, so C = 0 as well as U = 0.

```
Y at week 121 = 206
Y at week 154 = 186
n = 33 steps
ratio = 186/206 = 0.90291
delta = 1 − 0.90291^(1/33) = 0.003090 per week

half-life = ln 2 / (−ln(1 − 0.003090)) = 224.0 weeks = 4.31 years
tau = 1/delta = 323.6 weeks = 6.22 years
3 tau = 971 weeks = 18.7 years
```

The office's stored order has a four-and-a-third-year half-life, and its settling time to a new equilibrium is nearly nineteen years. Both facts matter in Chapter 26. The first says the manual is genuinely durable. The second says nobody now working in the office will see the equilibrium of any change made today.

Form E — eta

Window: the last 26 operating weeks of the series.

```
sum of F(U_t) over the window = 31.903 nats
Y-units created in the window = 26 (register count)
of these, attributable with a documented originating board decision = 16
unattributable, excluded = 10

eta = 16 / 31.903 = 0.5015 rules per nat
```

The 10 unattributable rules are themselves a finding: they were written by individual reviewers without a board decision behind them. They are real rules and they are in Y, but they were not produced by interaction, so they cannot be used to estimate the conversion of interaction into structure. **A high unattributable fraction is not a measurement nuisance. It is telling you that the board is not the thing making the order, which is a substantive claim about the office.**

Form F — cost

line	category	route	Y-units/week
1	Appeal reversals that strike a live rule	1	0.14

line	category	route	Y-units/week
2	Rules orphaned by departure (no qualified applier remains)	1	0.11
3	Rework: items re-reviewed after disposition	2	0.16
4	Contradiction maintenance	1	0.09
5	Attention tax at the declared conversion rate	2	0.13
	C_hat		**0.63**

Reconciliation: the cost implied by the observed series over the operating weeks is 0.6275 Y-units per week. Residual against the line-item tally: **0.0025**, against a declared tolerance of 0.05. The tally may be used.

Slack is not booked. Turnover is assigned by the standing rule: routine attrition sits in `delta`; the two departures that orphaned specific rules sit in line 2 of `C`.

Form G — Q, headroom, verdict

```
Y register count at survey date (week 312)  = 121
liveness correction (no application in 156 weeks) = −6
applier correction (orphaned, no qualified applier) = −2
Q = 113

eta ln N   = 0.5015 x 1.6094 = 0.8071
delta Q    = 0.003090 x 113  = 0.3492
C_hat      =                   0.6300
---------------------------------------
h          = 0.8071 − 0.3492 − 0.6300 = −0.1720
```

VERDICT: h < 0. Unsustainable in this model. No interaction geometry saves it.

This is Corollary 3.4 in its operative form. `C_hat = 0.63` exceeds `eta ln N − delta Q = 0.4579`, and therefore **no rearrangement of who sits on the board, in what order they speak, or how vigorously they disagree can prevent the manual from declining.** Even a perfectly orthogonal board running every week at `U = 1` would return `eta ln N − delta Q − C = −0.1720` rules per week.

The realised picture:

```
one-step = eta F_real − delta Q − C_hat
         = 0.5015 x 1.1498 − 0.3492 − 0.6300
         = 0.5766 − 0.9792
         = −0.4025 rules per week  =  −20.93 rules per year
```

```
realisation deficit = h − one-step = −0.1720 − (−0.4025) = 0.2305 rules/week
```

And the steady state, Theorem 3.5:

```
Q_inf = (eta F_real − C_hat) / delta = (0.5766 − 0.6300)/0.003090 = −17.3
```

`Q_inf` is negative. Reported as negative, per S24.27. Theorem 3.5 does not hide this failure mode: the equilibrium is negative, so the regime is unsustainable or the linear model must be replaced. Operationally it means the manual is being driven to the boundary of the model's validity — to zero live rules — and that the linear model says nothing useful below that boundary.

Projected trajectory at current parameters, from `Q = 113`:

horizon	Q
1 year	93.6
3 years	63.1
5 years	41.0
12.6 years	0

Why nobody noticed

The office lost 177 rules over the 312 weeks of the audit series — from 290 down to 113, a decline of 61%. At the observed rate of about 21 rules a year out of a stock of roughly 150, the quarter-to-quarter change is under 2%, which is inside the noise of the register and well inside the noise of anyone's impression. **The decline was invisible at every time scale on which the office actually looks at itself, and unmistakable on the one time scale nobody was using.**

That is the argument for the survey in a single sentence. The office's staff were not negligent and were not wrong about their own experience. They were reading a slow variable at a fast cadence.

The three measurements that would most change this conclusion (S24.29)

1. **The attention-tax conversion rate.** Line 5 of the cost tally, at 0.13 Y-units per week, is the only line resting on a declared conversion rather than a direct count. If the true rate is half what was declared, h rises to -0.107 — still negative, but the margin changes materially. Measuring displacement directly, by timing rule production against interrupted and uninterrupted weeks, would settle it.
2. **The unattributable fraction in the eta audit.** Ten of 26 new rules had no documented board origin. If half of those were in fact board-originated and simply undocumented, `eta` rises to about 0.658 and h becomes -0.0424 — close enough to zero to change the verdict from "unsustainable" to "maintenance only", which is a different chapter of Chapter 26 entirely.
3. **The liveness threshold.** Y uses 156 weeks of non-application as the liveness cutoff. At 104 weeks the count would be lower and `delta Q` smaller; at 260 weeks, higher. The cutoff was declared in advance and must not be moved after seeing the result, but its sensitivity should be reported, and it is: moving the cutoff by one year in either direction changes Q by about 5 rules and h by about 0.015.

None of these three would reverse the sign of h on its own. Two of them together would. **That is the honest statement of how firm this survey's verdict is, and it belongs in the survey rather than in a footnote.**

What the survey does not say

It does not say the office is badly run. It does not say the board is useless. It does not recommend anything.

It says six numbers and one sign. Everything that follows — whether the model can be trusted at all (Chapter 25) and what may be built if it can (Chapter 26) — is downstream of the fact that the sign is negative, and of the fact that we can now say exactly what it would take to change it.

Chapter 25 . Preregistration and Falsification

The survey of Chapter 24 produced six numbers. It did not produce a licence to use them.

Six numbers arranged into an equation are a hypothesis about how a real system behaves over time, and a hypothesis that has not been given a chance to fail is not a hypothesis. It is a decoration. This chapter is the procedure that gives the storage law a fair chance to be wrong about your organisation — and, when it is wrong, the procedure that makes you say so.

A framework that cannot fail cannot be trusted, and this one names exactly how it fails. That sentence is the whole chapter. Everything below is the machinery that makes it true rather than merely admirable.

The Law

The minimal estimation model (Paper E §5.2)

```
Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t

subject to      0 <= a <= 1        b >= 0

identification  a = 1 − delta      b = eta
```

This is the storage law written as something a computer can fit. `Y_t` is the order observable declared in Chapter 24. `F(U_t) = ln(1 + (N−1) U_t)` is the surplus input, computed from the interaction observable. `C_hat_t` is the cost observable, denominated in Y-units, and it enters with a **fixed coefficient of one** — it is a measured offset, not a free parameter. `eps_t` is the residual.

Two free parameters. That is all you get. The temptation to add a third — an intercept, a lag, a seasonal term, a second input — is the temptation to make the model unfalsifiable, and every one of those additions must be declared in advance or not made at all.

Why the constraints are the test, not a convenience

The box `0 <= a <= 1`, `b >= 0` is not regularisation. It is the empirical content of the claim.

`a = 1 − delta` with `delta` in `[0,1]` says the system is a **leaky accumulator**: it keeps some fraction of what it had and loses the rest. If the unconstrained best fit wants `a > 1`, the system is not accumulating with loss — it is compounding, and the storage law is the wrong model. If the unconstrained best fit wants `a < 0`, the system is oscillating period to period, which a stock does not do. If it wants `b < 0`, then interaction *destroys* order in your system, which is a coherent thing for a system to do and an incoherent thing for this model to describe.

Fit unconstrained first, and look at where the optimum lands. That is falsification rule one, and it is free.

The three falsification rules (Paper E §5.3, operational form)

Paper E §5.3 gives three rules under which the framework is rejected for a given system. Rendered as operating procedure, with the tolerances that you must declare in advance and may not move afterwards:

FR1 — THE ADMISSIBLE BOX. The unconstrained least-squares optimum must lie inside `0 <= a <= 1, b >= 0`. If it lies outside, the storage law is rejected for this system. No amount of constrained refitting rescues it; a constrained fit that sits on the boundary is reporting that the data wanted to leave the box.

FR2 — BASELINE DOMINANCE. The frozen model must beat **all four** required baselines on held-out data by a declared margin: persistence, direct-input, matched autoregressive, and the shuffled-input control. Failure against any one of the four is a failure of the rule. The recommended declared margins are a 10% reduction in held-out RMSE against each of the first three, and `p <= 0.05` against the shuffle distribution.

FR3 — INDEPENDENT-PARAMETER AGREEMENT. The fitted `a` and `b` must agree with the independently measured `delta` and `eta` from the survey, within a tolerance declared in advance. The recommended tolerance is +/-25% on each. Additionally, no observed one-period increase in `Y` may exceed the ceiling `eta ln N` by more than the declared conversion-lag allowance (Corollary 4.2).

Each rule catches a different disease, and this is why all three are required. FR1 catches a system that is not an accumulator at all. FR2 catches a model that adds nothing over inertia, trend, or the input's mere average level. FR3 catches a model that fits beautifully with parameters that are not your system's parameters — which is the signature of the order-proxy trap of Chapter 24 §24.15, and the only rule that reliably catches its most dangerous form.

Badges

POSTULATE (badged). Linearity. The storage law is linear in `Q` and in `S`. Real institutional stocks may saturate, may have thresholds, may exhibit hysteresis. The linear model is chosen for tractability and for the fixed-point guarantee, not derived from observation of organisations. A held-out failure may mean the storage law is wrong for your system, or it may mean the storage law is right and the linearity is wrong; the procedure below does not distinguish them, and you must not pretend it does.

POSTULATE (badged). Constant parameters. The model treats `delta`, `eta` and `N` as constant over the estimation window. In an organisation with turnover, reorganisation, or a change of leadership, they are not. Declare the window and check for a structural break; a fit that spans a regime change is fitting an average of two systems, neither of which exists.

POSTULATE (badged). Additive, independent, homoscedastic noise. `eps_t` is assumed to behave itself. Audit error on a register is usually neither independent nor identically distributed — audits taken by different people, in different weeks, under different pressure. Report the residual autocorrelation; if it is large, your held-out intervals are narrower than the truth.

The honesty clauses that apply specifically to estimation

- Paper A is pure mathematics and makes **no empirical or metaphysical claim.** Fitting the storage law to your organisation tests your model of your organisation. It does not test Paper A, and a failed fit is not a refutation of a theorem.
- The surplus functional as deployed is a **lower bound** on the ensemble's information. **We log the floor, not the ceiling.** A fitted `b` slightly above the audited `eta` is therefore not automatically an error; it may be the floor being conservative. A fitted `b` forty times the audited `eta` is not that.
- Paper D: surplus is **not** Shannon entropy or mutual information in general. Do not substitute an information-theoretic quantity for `F(U_t)` and call it the same model.
- Paper E proves **no thermodynamic negentropy law.** `F(u) + N_can(u) <= ln N` (Corollary 2.2).

The Engineering

25.1 What preregistration is for

Preregistration is the practice of writing down, in advance and in a form that cannot be silently edited, exactly what you will measure, exactly how you will model it, exactly what result would count as success, and exactly what result would count as failure.

It is not bureaucracy and it is not modesty theatre. It solves a specific, mechanical problem: with enough freedom in the analysis, any series can be made to support any conclusion, and the person with the most freedom is the person who most wants the conclusion. In organisational work that person is usually you, hired by someone with a preferred answer, working on a system whose staff will read your report.

The three freedoms that must be surrendered, in order of how much damage they do:

1. **Freedom over the outcome.** Choosing `Y` after seeing which candidate proxy moved.
2. **Freedom over the time scale.** Trying weekly, monthly and quarterly and reporting the one that worked. This one is especially seductive because, as §25.10 shows, the time scale genuinely matters — which is exactly why choosing it after the fact is fatal.
3. **Freedom over the specification.** Adding a lag, an intercept, a control, or a dropped outlier until the fit clears.

Preregistration surrenders all three in a dated document, before the estimation window's holdout period is examined.

25.2 The freeze

The freeze is the moment every choice becomes read-only. It has a date, a signature, and a copy held by someone who is not you.

What is frozen:

Frozen item	Source
The five operational choices	Chapter 24, Form A

Frozen item	Source
The repackaging audit results	Chapter 24, Form B
`N`, and the merge/drop decisions	Chapter 24, Form C
The estimation window and the train/holdout split	This chapter, §25.3
The time scale and its block alignment	§25.10
The model specification, including that `C_hat` has coefficient 1	§25.1
The four baselines and their exact definitions	§25.7
The three falsification rules and their numerical tolerances	The Law, above
The pre-declared handling of missing periods, outliers, and structural breaks	§25.4
The stopping rules and harm criteria for any intervention that follows	§25.13

What is *not* frozen, because it must not be: the holdout data. You may look at the training data as much as you like before the freeze — that is what training data is for. **The holdout is not examined until after the freeze, and it is examined once.**

If you must change something after the freeze, you may — and the change converts the entire analysis from confirmatory to exploratory, permanently, and must be labelled as such in every sentence that reports it. This is not a punishment. Exploratory analysis is how new hypotheses are found. It is simply not evidence.

25.3 Train, holdout, and the discipline of the split

Split the series by time, never at random. Random splits leak: a randomly held-out period sits between two training periods and its level is nearly determined by them.

Recommended split: the last 20-30% of the series, by calendar, is the holdout. Declare the exact boundary period in the preregistration.

Three rules that are violated constantly:

- **One look.** Compute the held-out numbers once. If you refit anything after looking, say so and relabel the analysis exploratory.
- **No holdout in the estimation of anything.** Including `delta` from the zero-input episode and `eta` from the attribution audit, if those episodes fall in the holdout window. If the zero-input episode is in the holdout, either move the split or accept that FR3's independent measurement is now dependent — and say which.
- **Report the holdout size.** A holdout of four blocks is not a holdout; it is an anecdote with a standard deviation. §25.10 shows what this costs.

25.4 Pre-declaring the awkward cases

Real series have gaps, spikes and regime changes. Decide in advance.

Periods with `U_t = 0`. Keep them. They are the most informative periods you have. Declare explicitly that `F(0) = ln(1) = 0`, which is exact and requires no special-casing.

Missing audits. Declare the rule: either carry forward and flag, or drop the pair and record the count. Never interpolate a stock; interpolation manufactures the smoothness you are trying to test for.

Outliers. Declare a rule that does not reference the residual, because a rule that drops large residuals is a rule that guarantees a good fit. The admissible rule is one that references the *record*: a period in which the register was migrated, recounted, or restated may be excluded, and the exclusion is documented at the freeze if known and reported as exploratory if discovered later.

Structural breaks. Declare the known ones — a merger, a new director, a legal change — and declare whether the window spans them. If it does, declare in advance whether you will fit separately on each side. Discovering a break after the fact and splitting on it is exploratory.

The ceiling check. Declare that any observed single-period increase exceeding `eta ln N` plus the conversion-lag allowance will be investigated before the analysis proceeds. This is FR3's arithmetic half and it is cheap.

25.5 The fitting procedure

Numbered, so it can be executed without judgement.

1. **Assemble the pair set.** For each period `t` in the training window for which both `t` and `t+1` are observed, form the row `(Y_t, F(U_t), C_hat_t, Y_{t+1})`.
2. **Move the cost to the left.** Define `z_t = Y_{t+1} + C_hat_t`. Because `C_hat` has a fixed coefficient of one, it is an offset and not a regressor.
3. **Fit unconstrained.** Solve the ordinary least squares problem `z_t ~ a Y_t + b F(U_t)` with no intercept. Record `(a_unc, b_unc)`. **This is FR1's test. Do it before the constrained fit and write the answer down.**
4. **Fit constrained.** Solve the same problem subject to `0 <= a <= 1`, `b >= 0` (any bounded linear least-squares routine will do). Record `(a, b)`.
5. **Translate.** `delta_fit = 1 − a`; `eta_fit = b`. Convert `delta_fit` to a half-life and a time constant so it is comparable with the survey's directly measured `delta`.
6. **Report the profile interval on a.** Refit `b` on a grid of fixed `a` values, record the sum of squared errors, and take the set of `a` within the declared confidence threshold. On slow stocks this interval is wide, and it is important to show that it is.
7. **Freeze the fitted coefficients.** They are now constants.
8. **Predict the holdout, once.** `Y_hat_{t+1} = a Y_t + b F(U_t) − C_hat_t`, one step ahead, using the *observed* `Y_t` at each step, not your own previous prediction. One-step-ahead is the honest evaluation for a stock; free-running multi-step prediction of a near-unit-root series is a test of your patience, not of your model.
9. **Compute held-out RMSE** for the model and for each of the four baselines.
10. **Run the shuffled-input control** (§25.8).
11. **Evaluate FR1, FR2, FR3 against the frozen tolerances.** Record each as PASS or FAIL. Do not average them. Do not weight them.
12. **Write the verdict sentence.** One sentence, of the form: *the storage law is / is not corroborated as a model of `Y` in this system at the preregistered time scale, and the rule it failed is ___.*

25.6 A note on what "beat the baseline" means

RMSE, one step ahead, on the holdout, in Y-units. Report the raw RMSE for every model, not only the skill scores, because a skill score of 40% means something completely different when the RMSE is 0.5 rules and when it is 50 rules.

Skill score convention used throughout this manual:

```
skill = 1 − RMSE_model / RMSE_baseline
```

Positive is better. A skill of 0.40 means the model's error is 60% of the baseline's.

25.7 The four required baselines, and the distinct failure each one catches

This is the part of the chapter that most often gets skimmed, and it is the part that does the work. **Each baseline is a specific accusation against your model. You are required to answer all four.**

Baseline 1 — PERSISTENCE. `Y_hat_{t+1} = Y_t`. No parameters.

The accusation: your model adds nothing over the observation that stocks are sticky.

What it catches: the commonest and most embarrassing failure, and the reason a stock is a treacherous thing to model. A quantity with a four-year half-life barely moves period to period, so "it will be what it was" is an excellent forecast. Any model of a slow stock that cannot beat persistence has learned nothing about the stock.

Note: persistence gets easier to beat as the period lengthens, because more happens between observations. A large skill against persistence at a coarse time scale is not impressive on its own; it must be read together with baselines 3 and 4.

Baseline 2 — DIRECT INPUT. `Y_hat_{t+1} = c_0 + c_1 F(U_t)`, with `c_0, c_1` fitted on the training data. No memory term.

The accusation: your Y is not a stock at all — it is a function of the input, and the memory term is decoration.

What it catches: **contemporaneous repackaging**, species 1 from Chapter 24 §24.15. If Y is an affine function of the interaction, then the input alone predicts it and the accumulator adds nothing. For a genuine stock this baseline is hopeless, because the input carries no information about the *level* — in the Worked Example below it has a held-out RMSE of 79.4 rules against the model's 1.85, a skill of 97.7%. **A direct-input baseline that performs well is not good news. It is the alarm for the order-proxy trap.**

Baseline 3 — MATCHED AUTOREGRESSIVE. `Y_hat_{t+1} = a_0 + a_1 Y_t`, with `a_0, a_1` fitted on the training data. Two free parameters, matching the storage model's two.

The accusation: your model is trend-following with extra steps; the input term earns nothing that memory plus drift did not already give you.

What it catches: the failure that matters most for institutional work, and by far the hardest of the four to beat. An AR(1) with an intercept can represent any exponentially-approaching trend. A declining institution is an exponentially-approaching trend. **If your storage model cannot beat**

a matched autoregression, then you have measured that the system is drifting, which you knew, and you have not measured that interaction is driving it, which is the entire claim.

The matching matters. A baseline with fewer free parameters than the model is not a fair comparison; the AR(1)-with-intercept has exactly two, as the storage model does.

Baseline 4 — THE SHUFFLED-INPUT CONTROL. Refit and re-evaluate the *same* storage model with the input series `F(U_t)` randomly permuted in time. Everything else — `Y_t`, `C_hat_t`, the constraints, the split — is unchanged. Repeat many times to build a null distribution of held-out RMSE, and locate your real fit within it.

The accusation: `b` is not fitting the input's dynamics; it is fitting the input's average level, and any series with the same distribution would have done as well.

What it catches: the failure that no other baseline can see. This is developed in §25.8.

This is the baseline most often omitted and most often decisive. It is omitted because it is the only one that requires writing a loop rather than a formula, and because — unlike the other three — it has no natural place in a table of "model versus alternatives". It is decisive because it is the only test that asks whether the *timing* of interaction matters, which is the only thing that distinguishes a driven accumulator from a bucket with a tap running at a constant average rate.

25.8 The shuffled-input control in detail

Procedure:

1. Take the training pair set. Permute the `F(U_t)` column, leaving `Y_t`, `C_hat_t` and `Y_{t+1}` in their original order.
2. Refit the constrained model on the permuted training data. You get `(a*, b*)`.
3. Permute the `F(U_t)` column of the *holdout* independently, and compute held-out RMSE with `(a*, b*)`.
4. Repeat at least 1,000 times. 2,000 is better and costs nothing.
5. Report: the mean and 5th percentile of the null RMSE distribution, the count of permutations achieving an RMSE at least as good as the real fit, and the resulting `p = (count + 1)/(n + 1)`.

Three details that are commonly got wrong:

- **Refit under the permutation.** If you keep the real `b` and only permute the input at prediction time, you have built a much weaker null — you are testing whether the input helps *given a coefficient fitted to the true alignment*, which is nearly circular.
- **Permute the holdout independently.** Otherwise the holdout retains the alignment your null is supposed to destroy.
- **Do not permute `C_hat`.** The control isolates the input term. Cost stays aligned, and stays with coefficient one, so its explanatory contribution is held constant between the real fit and the null. This is what makes the test specific.

What the shuffle can and cannot detect. It detects that the input's timing carries information. It does **not** detect that your `Y` is a filtered copy of your `U`, because a filtered input is genuinely predicted by the input, and the shuffle destroys that genuine prediction just as it would destroy a real one. The Worked Example below constructs exactly such a proxy: it passes the shuffle at `p = 0.0005`, beats persistence by 89.4%, beats the matched autoregression by 89.2%, and is

nothing but a rolling average of the input series. Only FR3 and the structural audit of Chapter 24 catch it.

The shuffle is the empirical audit; the repackaging tests are the structural audit; they are different audits and you need both.

Three measured failure modes, and the last two are not repairable by any surrogate. The paragraphs above state the second of them. Chapters 55 and 81 went on to measure the control's behaviour by simulation rather than argue about it, and the full list must be carried by every chapter that recommends this suite — this one above all, because this is where the suite is specified.

(1) It is anti-conservative as ordinarily run. Across 500 null trials against a nominal `alpha = 0.05`, the measured false-rejection rates were:

```
iid permutation      0.108      2.2x nominal
circular shift       0.120      2.4x nominal
block permutation    0.060      near nominal
```

A surrogate test run off the shelf roughly doubles the false-positive rate, and the block variant — the one most often skipped as fussy — is the only one that behaved. The `p <= 0.05` threshold in FR2 is therefore not a threshold you may take on trust. **Calibrate the control by simulation on your own study's design — your series length, your autocorrelation, your block structure — and preregister the calibrated threshold rather than the nominal one.** Chapter 55 makes this a requirement. In the same exercise the matched autoregressive baseline evaluated *alone* false-passed 35.8% of the time, which is the arithmetic behind S25.18: a suite is only a suite if all four are read together.

(2) It cannot see repackaging, as demonstrated above and confirmed by the Worked Example at `p = 0.0005`. The catch came from the structural ceiling test of Chapter 24 §24.15 Test 4 — a single-step jump **38.5 times** the maximum permitted by Corollary 4.2 — and from nothing else.

(3) It cannot see confounding by a shared slow driver. Where `U` and `Y` are both driven by a third slow variable — a funding cycle, a hiring wave, a seasonal market — a 4,000-trial simulation found the shuffled-input control **passing 98.35% of the time.** The same simulation reports the control working honestly in the clean case at 48.17% detection, which is what makes the 98.35% figure a measurement and not a grievance.

The consolidated rule, which supersedes any earlier reading of the shuffle as a general guarantee: the shuffle is a weak, mis-calibrated empirical audit. Repackaging is caught only by the structural tests; confounding and fabricated timing are caught only by design controls — blinding, channel separation, prospective sampling, and a yoked control. No surrogate catches either, and adding surrogates does not help. Those four controls are decisions about the instrument, not about the analysis, which means they must be written into Part 1 of the preregistration form and cannot be retrofitted once the series exists.

25.9 Reading the three rules together

```
FR1 fail -> the system is not a leaky accumulator. Stop. Different model needed.
FR2 fail -> the model has no demonstrated predictive advantage. You may not claim one.
FR3 fail -> the model fits with parameters that are not this system's. Suspect the proxy.
```

The combinations are informative:

FR1	FR2	FR3	Reading
pass	pass	pass	Corroborated at this scale. Proceed to Chapter 26 with a validated model.
pass	fail	pass	Parameters corroborated; predictive advantage not demonstrated. **Accounting frame only** — see §25.11.
pass	pass	fail	The classic order-proxy trap. Return to Chapter 24 Form B.
pass	fail	fail	Uninformative. The series is too short, too noisy, or the proxy is wrong.
fail	any	any	Rejected. The storage law is not the right model for this system.

Note the second row carefully, because it is the outcome the Worked Example produces and it is common in real institutional data. **Corroborated parameters with no demonstrated predictive advantage is a real and useful state, and it licenses much less than people want it to.**

25.10 Choosing the time scale before you look

Chapter 24 §24.5 gave three constraints on the period. Here is the fourth fact, which is the reason the choice must be frozen: **the time scale materially changes the result, and it changes it in a direction you can anticipate.**

The structure is this. Audit error on a *level* does not accumulate — if you count the register twice, six months apart instead of one week apart, the counting error at each endpoint is the same size. But the accumulated input between those endpoints has grown with the interval. So coarsening the time scale raises the signal-to-noise ratio of the `b F(U)` term. It also, linearly, destroys your degrees of freedom.

There is therefore an optimum, and the optimum depends on quantities you can estimate before looking at the holdout:

```
signal per block  ~  B x eta x mean F(U)
noise per block   ~  sqrt(2) x (audit error SD)        (does not grow with B)
blocks available  =  T / B
```

Read the first line carefully: it is `mean F(U)`, and the distinction is not decoration. When you coarsen the time scale you are aggregating `U` over `B` fine periods, and there are two ways to carry

the input into the block — average the U values and take F once, or take F of each fine period and average the results. They are different numbers. F is strictly concave (Paper A, derived result 3.3), so Jensen's inequality fixes the direction:

```
mean of F(U_t) <= F( mean of U_t )      with equality only if U is constant within the block
```

Aggregate F, never aggregate U. Taking `F(mean U)` inflates the input term, and it inflates it more the coarser the block — which means a preregistration that aggregates the wrong way will show the model improving at every increase in B for reasons that have nothing to do with signal-to-noise, and the engineer will read a Jensen artefact as a discovery about the time scale. The gap is worth having a name and a number: it is the **concavity penalty**, and in the Chapter 26 worked office it comes to 0.0761 nats against an `F(mean U)` of 1.2262 — **6.2% of the surplus**, lost purely to the variability of the geometry. Declare the aggregation rule in Part 1 of the form alongside the aggregator itself, and report the penalty with the fit. The same rule governs the feasibility test in Chapter 26 §26.7 and the headroom arithmetic in Chapter 24 §24.14.

Choose B by writing the trade-off down and picking a value, with a stated reason, before the freeze. Legitimate reasons include: B must exceed the conversion lag; B must give at least 8 holdout blocks; B must not push `delta` per block above 0.3.

Illegitimate reasons include: `B = 13` gave the best result.

In the Worked Example, the answer at `B = 13` is materially better than at `B = 8` on every metric. It is also based on four holdout blocks, and it was not the preregistered choice, so it is reported as exploratory and it licenses nothing. That is what the discipline costs, and paying it is the difference between a survey and an advertisement.

25.11 What a passing fit does and does not license

A pass licenses: using the fitted a and b as corroborated estimates of `delta` and `eta` for the system, at the fitted time scale, over the fitted regime; computing h from them; computing a maximum admissible increment `r <= h`; and stating a one-step prediction with an interval derived from the held-out residuals.

A pass does not license:

- **Multi-period forecasts.** The model was evaluated one step ahead. A ten-step forecast of a near-unit-root series has an error that grows without any of the discipline above having tested it.
- **Causal claims.** The shuffle test establishes that the timing of interaction carries information about subsequent order. It does not establish that intervening on interaction will produce order. That is a different claim requiring a different design — Chapter 26 §26.11's pilot, with its own preregistration.
- **Transfer to another system.** `delta`, `eta`, `N` and `C` are the parameters of *this* system. The scale-free form of the law is a claim about model form, not about parameter values. Hyper-recursive symmetry says the same equation applies at desk, room, team and institution; it says nothing about the numbers being the same.
- **Transfer to another proxy.** A fit validated on live precedent rules says nothing about reviewer competence, cycle time, or public satisfaction.
- **Extrapolation past the ceiling.** `Q_t <= max{Q_0, eta ln N / delta}` (Theorem 3.3). A prediction above the uniform bound is a prediction the model has already forbidden.

25.12 Goodhart's Law, Campbell's Law, and why preregistration is a corruption-control device

Goodhart's Law: when a measure becomes a target, it ceases to be a good measure. Campbell's Law makes the sociological version explicit: the more a quantitative indicator is used for social decision-making, the more it will distort and corrupt the processes it was meant to monitor.

Both apply to Y, immediately and mechanically, from the moment anyone in the organisation learns that Y is being counted.

The mechanism in this framework's own terms is precise. Y is a proxy for Q. Goodharting is the process by which the correlation between Y and Q is destroyed by effort directed at Y. A review office that learns its live-rule count is the metric can raise the count without raising the stored order at all: split existing rules, write rules nobody applies, restate one rule as three. The proxy goes up, the order does not move, and every number downstream — `delta`, `eta`, h — becomes a measurement of the gaming rather than of the system.

Five defences, in descending order of effectiveness:

1. **Keep Y out of the reporting line.** The strongest defence by a wide margin. A proxy that is not used to evaluate anyone is not worth gaming. Where Y must be reported, report h and the verdict, not Y.
2. **The liveness correction as standing practice.** Chapter 24 §24.12. Gamed units are created and never applied, so a liveness threshold removes most of them automatically and does so without anyone having to accuse anyone.
3. **Preregister the proxy definition, including its corrections.** Goodharting works by exploiting the *edges* of a definition. A definition frozen and dated before the incentive existed cannot be quietly widened.
4. Rotate the proxy on a schedule known to the surveyor and not to the surveyed. Expensive, and it breaks series continuity, so use it only where the incentive is strong.
5. **Measure at least one proxy the organisation cannot influence.** In the Worked Example, appeal reversal counts come from an external appeals board. A proxy the organisation does not control is a check on the ones it does.

Preregistration is not only an epistemic device. It is the specific administrative countermeasure to Campbell's Law, because it fixes the definition of the measure at a time when nobody yet has a reason to distort it.

25.13 Preregistering an organisational intervention honestly

Everything above concerns modelling an existing system. If the preregistration also covers an intervention — and Chapter 26's pilot requires that it does — five more sections are mandatory.

1. The prediction, in numbers, with a direction and a magnitude. Not "we expect improvement". The form is: *Y will increase by at least `r` per period, measured at the preregistered time scale, over at least k periods, beginning no earlier than period `m` after the change.* The lag `m` must be stated because a conversion lag exists and testing before it has elapsed is a test you have designed to fail.

2. The falsification criterion for the intervention, distinct from FR1-FR3. What observation would make you say the intervention did not work? State it as a number. The standard form: *if*

the mean one-period change over periods `m` to `m + k` is below `r_min`, the intervention is judged ineffective and will be withdrawn.

Preregister a second falsification alongside it, because it is free. Hold `D`, `p_bar`, `N` and the budget constant and let u decay geometrically at rate `lambda`; the doctrine-to-data ratio of Chapter 48 crosses 1 at

```
t* = ln(u*/u_0) / ln(1 − lambda)          periods
```

`t*` contains no free parameters. Every term is already on the survey form — the measured `u_0`, the crossing level u∗, the measured decay `lambda` — so there is nothing to fit, nothing to tune, and no researcher degree of freedom left to exercise. It yields a date. In the worked case of Chapter 48 the date falls at about 8.4 quarters. **A prediction that names a calendar date and costs nothing to make is the cheapest decisive falsification test available anywhere in this part of the manual**, and there is no defensible reason to omit it from a preregistration that already commits to a series.

And declare the unit of analysis, because `r` does not aggregate. Distortion has a sign, and opposite-signed distortion cancels under aggregation without leaving a trace. In the two-site example of Chapter 48, one over-framed site and one under-framed site combined into an aggregate `D/P` that **overstated bankable input by 4.26x**, while each site analysed alone was diagnosed correctly. **Never aggregate `r` across sites, divisions or cohorts.** Preregister the per-site analysis and report a vector of `r` values; a pooled `r` near 1 is more often two opposite pathologies than one healthy system, and by construction the pooling destroys exactly the evidence that would have told you which.

3. The stopping rules. Two kinds, both required. - *Futility:* the condition under which you stop because it is not working. - *Harm:* the condition under which you stop immediately regardless of how well `Y` is doing. Harm criteria must include at least one measure that is not `Y`, because an intervention that raises the proxy while damaging the system will not stop itself.

4. The consent and declination clause. CANON §5.4 is binding on this manual: no intervention that removes a person's ability to decline is in scope. The preregistration must therefore state, for each affected group: what is being changed, how they were informed, and by what mechanism an individual may decline without penalty. If the honest answer to the last is "they may not", the intervention is out of scope and the preregistration does not proceed. This is not a formality appended to the science; it is a constraint on which interventions exist at all, and Chapter 26 §26.16 treats it as a design gate.

5. The reversibility statement. What is the exact procedure for returning the system to its prior state, who is authorised to execute it, how long does it take, and what is not recoverable? Chapter 26 makes reversibility a hard requirement; the preregistration is where it is written down while it is still cheap to be honest about.

A sixth item is optional and strongly recommended: **the pre-mortem.** Before the freeze, write the paragraph you would write in a year explaining why the intervention failed. It costs twenty minutes and it routinely surfaces the harm criterion you had not thought of.

The Specification

S25.1 A preregistration document shall exist, be dated, be signed, and be held in a copy by a person who is not the analyst.

S25.2 The freeze shall occur before any examination of holdout data.

S25.3 The model specification shall be exactly `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t` with `0 <= a <= 1` and `b >= 0`, unless a deviation is declared at the freeze with its reason.

S25.4 `C_hat_t` shall enter with a fixed coefficient of one. Estimating a coefficient on cost is a declared deviation and changes the identification.

S25.5 No intercept shall be added unless declared at the freeze with a stated interpretation.

S25.6 The train/holdout split shall be by calendar time, never at random, and the boundary period shall be named.

S25.7 The holdout shall contain at least 8 estimation blocks. If it cannot, the analysis shall be labelled exploratory throughout.

S25.8 The unconstrained fit shall be computed and reported before the constrained fit.

S25.9 FR1 shall be evaluated on the unconstrained optimum, not on the constrained one.

S25.10 All four baselines shall be computed: persistence, direct-input, matched autoregressive, and shuffled-input.

S25.11 The matched autoregressive baseline shall have the same number of free parameters as the storage model.

S25.12 The shuffled-input control shall refit the model under each permutation, permute the holdout independently, and use at least 1,000 permutations.

S25.13 The shuffled-input control shall not permute `C_hat`.

S25.14 Raw held-out RMSE shall be reported for the model and for all four baselines, in `Y`-units, alongside any skill scores.

S25.15 The margins for FR2 shall be declared at the freeze. The recommended defaults are 10% RMSE reduction against each of the first three baselines and `p <= 0.05` against the shuffle.

S25.16 The tolerance for FR3 shall be declared at the freeze. The recommended default is +/-25% on each of `delta` and `eta`.

S25.17 FR1, FR2 and FR3 shall each be reported as PASS or FAIL individually. They shall not be averaged, weighted, or combined into a score.

S25.18 A failure of any one of the four baselines is a failure of FR2.

S25.19 The time scale shall be chosen and justified before the freeze. Results at other time scales may be reported only as exploratory and only with that label attached in the same sentence.

S25.20 The fitted a shall be reported with a profile interval, and the corresponding interval on the half-life shall be stated in calendar units.

S25.21 Held-out prediction shall be one step ahead using observed `Y_t`, not free-running.

S25.22 Residual autocorrelation shall be reported.

S25.23 Periods with `U_t = 0` shall be retained, with `F(0) = 0`.

S25.24 The outlier rule shall not reference residuals.

S25.25 Known structural breaks shall be declared at the freeze, with the handling rule.

S25.26 The verdict shall be written as one sentence naming the rule that failed, if any.

S25.27 A verdict of "parameters corroborated, predictive advantage not demonstrated" shall be reported in those words and shall not be paraphrased upward.

S25.28 Where the preregistration covers an intervention, it shall state the numeric prediction, the falsification criterion, both stopping rules, the consent and declination mechanism, and the reversibility procedure.

S25.29 At least one harm criterion shall be a measure other than Y.

S25.30 If no mechanism exists by which an affected individual may decline without penalty, the preregistration shall not proceed and the reason shall be recorded.

The Preregistration Form

Printable. Sign at the bottom. Once signed, read-only.

```
================================================================
PREREGISTRATION                              Freeze date: ______
System: ........................  Analyst: ..............
Independent holder of this copy: ............................
================================================================

---- PART 1  CARRIED FORWARD FROM THE SURVEY ------------------
1.1  Survey document reference and date .....................
1.2  Y = ...................................................
1.3  Y units ................................................
1.4  U aggregator ...........................................
1.5  N = _____   ln N = _____
1.6  delta (measured, method _____) = _____ per period
1.7  eta (attribution audit) = _____ Y-units per nat
1.8  C_hat = _____ Y-units per period
1.9  Q at survey = _____
1.10 Repackaging tests 1-4 all PASS?  [ ] Y  [ ] N (if N, stop)

---- PART 2  ESTIMATION DESIGN --------------------------------
2.1  Series: periods _____ to _____  (T = _____)
2.2  Block length B = _____   Alignment: ....................
2.3  JUSTIFICATION for B (must not reference results):
     conversion lag = _____ ; blocks available = _____ ;
     holdout blocks = _____ ; delta per block = _____
     ..........................................................
2.4  Train: blocks _____ to _____      (n = _____)
```

```
2.5  Holdout: blocks ____ to ____    (n = ____)
2.6  Model:  Y_{k+1} = a Y_k + b F(U_k) - C_hat_k + eps_k
     constraints 0 <= a <= 1, b >= 0 ; C coefficient fixed at 1
2.7  Deviations from the minimal model (if any) and reasons:
     ....................................................
2.8  Missing-period rule ..................................
2.9  Outlier rule (must not reference residuals) ..............
2.10 Known structural breaks and handling ...................

---- PART 3  BASELINES ----------------------------------------
3.1  Persistence:      Y_hat = Y_k
3.2  Direct input:     Y_hat = c0 + c1 F(U_k), fitted on train
3.3  Matched AR(1):    Y_hat = a0 + a1 Y_k, fitted on train
3.4  Shuffled input:   refit under permutation, n_perm = ____
3.5  Metric: one-step-ahead held-out RMSE in Y-units

---- PART 4  FALSIFICATION RULES AND TOLERANCES ---------------
4.1  FR1  unconstrained optimum inside 0<=a<=1, b>=0
4.2  FR2  margin vs persistence ............. ____ % RMSE
          margin vs direct input ............ ____ % RMSE
          margin vs matched AR .............. ____ % RMSE
          shuffle p threshold ............... ____
4.3  FR3  tolerance on delta ....... +/- ____ %
          tolerance on eta ......... +/- ____ %
          conversion-lag allowance for the ceiling check ____
4.4  Any single FR failure = the framework is not corroborated
     for this system at this time scale.       [ ] initialled

---- PART 5  INTERVENTION (complete only if one is planned) ---
5.1  Change to be made ....................................
5.2  Parameter(s) targeted  [ ]eta [ ]C [ ]u [ ]N [ ]shed Q
5.3  PREDICTION: Y will rise by at least r = ____ per period,
     over k = ____ periods, starting m = ____ periods after
     the change (m >= conversion lag).
5.4  Headroom check: r <= h = ____             [ ] confirmed
5.5  FALSIFICATION: if mean one-period change over periods
     m..m+k is below ____, the intervention is ineffective
     and will be withdrawn.
5.6  STOPPING RULE, futility ..............................
5.7  STOPPING RULE, harm (must include a non-Y measure) .......
     ....................................................
5.8  CONSENT: who is affected ..............................
     how informed ..........................................
     how an individual may decline WITHOUT PENALTY .............
     [ ] a declination mechanism exists  (required; if not, stop)
5.9  REVERSIBILITY: procedure .............................
     authorised to execute .................................
     time to reverse .......................................
```

```
     what is NOT recoverable ..................................
5.10 PRE-MORTEM (why this failed, written in advance):
     ..........................................................
     ..........................................................

---- PART 6  SIGNATURES ------------------------------------
Analyst ......................  Date ..........
Independent holder ............  Date ..........
Sponsor / system owner .......  Date ..........
============================================================
```

The Failure Modes

F25.1 — THE MOVED GOALPOST. *Description.* A tolerance or margin is adjusted after the held-out numbers are seen. *Damages.* The entire evidential status of the analysis. *Signature.* A reported margin that is not a round number, or that differs from the preregistration by a small amount in exactly the direction needed. *Detection.* Compare the report to the frozen document held by the independent holder. This is what the independent holder is for. *Fix.* Relabel as exploratory, in every sentence, and rerun the confirmatory analysis on a fresh holdout in the next cycle.

F25.2 — THE OMITTED SHUFFLE. *Description.* Three baselines are reported and the shuffled-input control is not. *Damages.* The claim that interaction is doing anything. *Signature.* A results table with exactly three baseline rows. *Detection.* Trivial, and it is the first thing to look for in any report using this framework. *Fix.* Run it. It takes a loop and a minute of compute. Expect it to change the conclusion in a meaningful fraction of cases.

F25.3 — THE UNMATCHED AUTOREGRESSION. *Description.* The AR baseline is given fewer free parameters than the storage model, usually by dropping the intercept. *Damages.* FR2, in the model's favour. *Signature.* An AR baseline described as "AR(1)" with one coefficient reported. *Detection.* Count the parameters on both sides. *Fix.* Restore the intercept and rerun. Expect the model's apparent advantage to shrink substantially; on slow stocks it usually does.

F25.4 — MULTI-STEP FLATTERY. *Description.* Held-out performance is computed on free-running multi-step forecasts, where persistence and AR degrade fast and the model looks heroic. *Damages.* FR2, in the model's favour, sometimes by an order of magnitude. *Signature.* Skill scores above 0.9 against persistence on a stock with a long half-life. *Detection.* Ask which `Y_t` was used at each step: the observed one, or the model's own previous prediction. *Fix.* One-step-ahead with observed `Y_t`. If multi-step performance is genuinely of interest, preregister it separately, with its own baselines evaluated the same way.

F25.5 — TIME-SCALE SHOPPING. *Description.* Several block lengths are tried and the best is reported as the analysis. *Damages.* Everything; this is p-hacking with an unusually large search space, because the block length also changes the effective sample size. *Signature.* A block length with no stated justification, or a justification that could equally have supported two other values. *Detection.* Check Part 2.3 of the preregistration for a justification that does not reference results. *Fix.* Report the preregistered scale as confirmatory and all others as exploratory, in the same table, with the label in the row.

F25.6 — THE COST COEFFICIENT. *Description.* `C_hat` is entered as a regressor with a free coefficient because the fit is better. *Damages.* The identification. With a free coefficient on cost, b is no longer `eta` and h can no longer be computed from the fit. *Signature.* Three fitted coefficients reported instead of two. *Detection.* Count them. *Fix.* Fix the coefficient at one. If the fit is materially worse, that is information: your cost observable is mis-denominated, and the repair is in Chapter 24 §24.6, not in the regression.

F25.7 — THE RESIDUAL-BASED OUTLIER RULE. *Description.* Periods with large residuals are excluded as "data quality issues". *Damages.* Guarantees a good fit and destroys FR2. *Signature.* An exclusion list whose members are all large residuals. *Detection.* Ask what the exclusions have in common in the *record*. If the answer is only "the model predicted them badly", this is the fault. *Fix.* Restore them. Investigate them individually against the record; genuine restatements can be excluded with a documented, record-based reason.

F25.8 — THE UNLABELLED EXPLORATORY RESULT. *Description.* An analysis run after the freeze, or after a deviation, is reported alongside confirmatory results without a label. *Damages.* The reader's ability to weight the evidence. *Signature.* A results section in which every number has equal typographic status. *Detection.* Cross-reference every reported number against the frozen document. *Fix.* Label in the sentence, not in a footnote. The convention that works: the word "exploratory" appears in the same sentence as the number.

F25.9 — THE ETHICS APPENDIX. *Description.* Consent, declination, stopping and reversibility are written after the design is settled, as a compliance step. *Damages.* The design itself, invisibly, because an intervention designed without a declination mechanism usually cannot have one retrofitted. *Signature.* Part 5.8 of the form completed last, in different handwriting. *Detection.* Ask what the design would have been if declination had been a constraint from the start. If the answer is "different", the fault is present. *Fix.* Redesign. **CANON §5.4 is not a review gate; it is a constraint on the space of admissible designs, and it belongs at the front of the process.**

F25.10 — SUCCESS DECLARED FROM AN ACCOUNTING FRAME. *Description.* FR2 fails, FR3 passes, and the report says the model "was validated". *Damages.* Everything downstream, because Chapter 26's pilot rules depend on which of these happened. *Signature.* The phrase "broadly consistent with" doing load-bearing work. *Detection.* Look for the individual PASS/FAIL lines. If they are not individually reported, S25.17 has been violated. *Fix.* Report the verdict in the prescribed words: *parameters corroborated, predictive advantage not demonstrated.* Then apply §25.11's licence list, which permits computing h and forbids claiming a forecast.

Worked Example

The system is the **Meridian County Development Review Office** of Chapter 24. **Badge: constructed composite; the arithmetic is real, computed and reproducible; the organisation is invented and the audit series is a simulation with declared generating parameters. It must not be cited as evidence about any real office.**

Carried forward from the survey

```
Y            = live precedent rules in the standing review manual
N            = 5,  ln N = 1.6094
```

```
delta        = 0.003090 per week   (Method D1, weeks 121–154 freeze)
eta          = 0.5015 rules per nat (attribution audit, 26 weeks)
C_hat        = 0.63 rules per week
Q at survey  = 113
h            = −0.1720            VERDICT: h < 0
Repackaging tests 1–4: all PASS
```

The frozen design

```
Series:           weeks 1 to 312
Block length:     B = 8 (two-month blocks)
Justification:   conversion lag measured at up to 6 weeks, so B must exceed 6;
            B = 8 leaves 8 holdout blocks at the declared split, meeting S25.7;
               delta per block = 1 − (1−0.003090)^8 = 0.0244, well under 0.3.
Split:            train = blocks beginning on or before week 240 (n = 29)
                  holdout = blocks beginning after week 240 (n = 8)
Model:           Y_{k+1} = a Y_k + b F(U_k) − C_hat_k + eps_k,  0<=a<=1, b>=0
FR2 margins:      10% RMSE against persistence, direct input and matched AR;
                  shuffle p <= 0.05
FR3 tolerance:   +/− 25% on delta and on eta; conversion-lag allowance 6 weeks
Permutations:     2000
```

Frozen and signed before the holdout was examined.

The fit at the preregistered scale, B = 8

```
Unconstrained optimum:   a = 0.97439    b = 0.5859
Constrained optimum:    a = 0.97439   b = 0.5859    (interior; constraints not binding)

delta implied per block = 1 − 0.97439 = 0.02561
delta implied per week  = 1 − 0.97439^(1/8) = 0.003237
half-life implied       = 4.11 years
```

Held-out performance, one step ahead, 8 blocks

Model	held-out RMSE (rules)	skill vs this baseline
Storage model (frozen)	**1.852**	—
Persistence	3.102	+0.403
Direct input	79.397	+0.977
Matched AR(1) with intercept	1.906	+0.028
Shuffled input, mean of 2000	3.400	—
Shuffled input, 5th percentile	2.657	—
Shuffled input, permutations beating the fit	0 of 2000	p = 0.0005

Evaluating the three rules

FR1 — ADMISSIBLE BOX. The unconstrained optimum is `a = 0.97439`, `b = 0.5859`. Both lie strictly inside the box; the constraints are not binding. **PASS.** The office's manual behaves like a leaky accumulator: it retains most of what it has and gains in proportion to a non-negative multiple of the surplus input.

FR2 — BASELINE DOMINANCE.

Baseline	required margin	achieved	result
Persistence	10% RMSE reduction	40.3%	PASS
Direct input	10% RMSE reduction	97.7%	PASS
Matched AR(1)	10% RMSE reduction	**2.8%**	**FAIL**
Shuffled input	`p <= 0.05`	`p = 0.0005`	PASS

FR2 FAILS. The storage model beats a matched autoregression by 2.8% on held-out RMSE, against a preregistered requirement of 10%. Under S25.18, failure against any one of the four baselines is a failure of the rule, and the rule is not averaged.

FR3 — INDEPENDENT-PARAMETER AGREEMENT.

Parameter	survey (independent)	fit	difference	tolerance	result
`delta` per week	0.003090	0.003237	+4.8%	+/-25%	PASS
`eta` (= b)	0.5015	0.5859	+16.8%	+/-25%	PASS
Ceiling check	max weekly rise 2 rules vs `eta ln N = 0.8071` over a 6-week lag allowance	—	—	—	PASS

FR3 PASSES. Note the direction of the `eta` discrepancy: the fit is 16.8% *higher* than the attribution audit. That is the expected direction, and it is consistent with two things at once — the audit's deliberate strictness (unattributable units excluded, Chapter 24 Form E) and the honesty clause that the surplus functional as deployed is a lower bound. **We log the floor, not the ceiling**, so a fitted conversion slightly above the audited one is not an anomaly to be explained away.

The verdict sentence (S25.26)

The storage law is not corroborated as a predictive model of live precedent rules in the Meridian County Development Review Office at the preregistered eight-week time scale; it failed FR2 against the matched autoregressive baseline at a margin of 2.8% against a required 10%, while passing FR1 and FR3.

Row two of the table in §25.9: **parameters corroborated, predictive advantage not demonstrated.**

Exploratory results at other time scales

Reported for completeness. **Every row except `B = 8` is exploratory and licenses nothing.**

B	status	n train	n hold	a	b	delta/week	half-life (yr)	RMSE model	RMSE pers	RMSE direct	RMSE AR	skill vs pers	skill vs AR	shuffle p
1	exploratory	140	71	0.9970	7.5389	0.00294	4.54	1.372	1.439	80.015	1.387	+0.046	+0.011	0.0015
2	exploratory	120	35	0.9942	4.5299	0.00284	4.61	1.513	1.558	79.920	1.386	+0.029	-0.091	0.0030
4	exploratory	59	17	0.9876	8.5674	0.00309	4.30	1.541	1.863	79.882	1.450	+0.173	-0.063	0.0005
8	**confirmatory**	**29**	**8**	**0.9743**	**9.5859**	**0.00324**	**4.11**	**1.852**	**3.102**	**79.397**	**1.906**	**+0.403**	**+0.028**	**0.0005**
13	exploratory	17	4	0.9525	3.6461	0.00373	3.56	1.209	4.637	76.914	1.980	+0.739	+0.389	0.0155

Four observations, all of which are why S25.19 exists.

1. **The shuffled-input control passes at every scale**, at p between 0.0005 and 0.0155. The timing of interaction carries real information about the subsequent rule count at all five scales tested. This is the strongest single positive result in the analysis.
2. The matched autoregression is beaten at only two scales, and at `B = 2` and `B = 4` it beats the storage model outright. A declining stock is an exponentially-approaching trend, and an AR(1) with an intercept represents that trend exactly. This is the hard baseline, and it is hard for a structural reason rather than an accidental one.
3. **`B = 13` would have passed everything**, at +73.9% against persistence and +38.9% against the matched autoregression. It was not the preregistered scale and it rests on four holdout blocks. Reporting it as the result would have been the exact fault named in F25.5. It is reported here as exploratory, and the correct action it licenses is not a claim but a plan: extend the series and preregister `B = 13` for the next cycle, with the holdout requirement met.
4. **The implied `delta` is stable across scales** — 0.00289 to 0.00373 per week, against a directly measured 0.00309 — which is a genuine and reassuring consistency, and is FR3 doing its job independently of FR2's outcome.

The two repackaging demonstrations

The office was, at one point in the survey, offered two alternative order observables. Both were tested. Both are instructive.

Candidate PD-1 — "items dispositioned per session" (contemporaneous repackaging). Proposed because the minutes already contain it. It shares source fields with `U` and fails Chapter 24's Test 1 immediately, so it never reached the estimation stage — but it was fitted anyway, as a demonstration.

Model	held-out RMSE
Storage model	3.079
Persistence	3.314

Model	held-out RMSE
Direct input	**2.443**
Matched AR(1)	2.657

Fitted a = 0.7106, b = 13.759; shuffle p = 0.2278.

The direct-input baseline beats the storage model outright — a skill of -0.261. That is baseline 2 performing its designed function: when the input alone predicts your order observable better than an accumulator driven by the input, your order observable is a function of the input. The shuffle also fails, at p = 0.2278. PD-1 is caught twice over and would have been caught for free by Chapter 24's Test 1.

Candidate PD-2 — "the board disagreement index", a rolling weekly index published by the board clerk from the minutes (filtered repackaging). This one is dangerous.

Model	held-out RMSE	skill
Storage model	**0.833**	—
Persistence	7.870	+0.894
Direct input	12.940	+0.936
Matched AR(1)	7.720	+0.892
Shuffled input, mean of 2000	7.842	p = 0.0005

PD-2 passes all four baselines, decisively, on every margin. It beats persistence by 89.4%, the matched autoregression by 89.2%, direct input by 93.6%, and the shuffled-input control at p = 0.0005. Against the actual proxy's marginal 2.8% against the autoregression, PD-2 looks like the far better model of the office.

PD-2 is a rolling average of the input. It stores nothing. There is no institution behind it.

What catches it:

Test	Result
FR3, `eta`	fitted b = 20.205 against an audited eta = 0.5015 — a factor of **40**, far outside the +/-25% tolerance. **FAIL.**
FR3, `delta`	fitted a = 0.7038, implying a half-life of **1.97 weeks** against a directly measured **224 weeks** — a factor of **114**. **FAIL.**
Chapter 24 Test 3, zero-input persistence	across the 34-week freeze, PD-2 fell 82.1 -> 57.9 -> 28.5 -> 6.7 -> 1.6 -> 0.1, and then returned to **31.6 in the single week the board resumed**. The true proxy went 206 -> 186 and continued from 186. **FAIL.**

Test	Result
Chapter 24 Test 4, the ceiling	that single-week jump is 31.1 units. The largest admissible one-step increment is `eta ln N = 0.8071`. The observed jump is **38.5 times the theoretical maximum**. **FAIL.**

A stock cannot teleport. The ceiling test is arithmetic, requires no statistics, takes one line, and is the single most reliable detector of the order-proxy trap in this manual.

And the lesson that governs the architecture of these three chapters: **the shuffled-input control is the empirical audit and it cannot see repackaging; the structural tests of Chapter 24 and FR3 are the definitional audit and they can. Both are required, and they are required in that order, because the structural audit is free and the empirical one is not.**

What the office is licensed to do

Under §25.11, with FR1 and FR3 passed and FR2 failed:

Licensed. Treat `delta = 0.003090`, `eta = 0.5015`, `N = 5`, `C_hat = 0.63` and `Q = 113` as corroborated parameters of the system at the eight-week scale. Compute `h = −0.1720`. Conclude that the arrangement is unsustainable in this model. Use `h` as a budget constraint on any proposed increment.

Not licensed. Claim that the storage law forecasts the manual's rule count better than a trend line does. Claim that raising interaction will raise the rule count, on the strength of this analysis — the shuffle establishes that the timing of interaction carries information, not that intervening on it produces order. Forecast beyond one step. Extend any of it to cycle time, staff competence, or applicant satisfaction.

Required next. Chapter 26 proceeds on the accounting frame only: a capped, reversible pilot, preregistered under Part 5 of the form above, whose own falsification criterion is stated in numbers before it begins — and which, if it works, supplies exactly the causal evidence this analysis is not entitled to assume.

Chapter 26 . Intervention Design

Chapter 24 measured the system. Chapter 25 gave the model a fair chance to be wrong about it. This chapter is where you are finally permitted to touch something — and it begins by telling you, in one number, whether you are permitted to touch it at all.

The number is h. **Measure h before you promise anything, before you design anything, and before you accept the brief.** Every other decision in this chapter is downstream of its sign.

The Law

The gauge

```
h_t := eta ln N − delta Q_t − C_t                    (Paper E, section 6.4)

h_t <  0    unsustainable in this model. No interaction geometry saves it.
h_t =  0    maintenance only. Hold; do not build.
h_t >  0   growth available; any requested increment must satisfy r_t <= h_t.
```

Read it as a budget. `eta ln N` is the income at the geometric ceiling. `delta Q_t` is the mortgage on what is already stored. `C_t` is operating expense. `h_t` is what remains.

The four theorems that make this a decision procedure

Maximum one-step increment (Corollary 4.2). `r_max = eta ln N − delta Q_t − C_t = h_t`. **No policy based on interaction geometry alone can guarantee more.** This is not a rule of thumb; it is the extremal statement, and `r_max` is attained only at `U_t = 1`, that is, only if every category is mutually orthogonal to every other.

Target increment feasibility (Proposition 4.1). `Q_{t+1} − Q_t >= r_t` if and only if

```
S_t >= (delta Q_t + C_t + r_t) / eta
```

and this is feasible only if `delta Q_t + C_t + r_t <= eta ln N` — which is exactly `r_t <= h_t`. The proposition therefore gives you two distinct checks, and both must pass: is the increment inside the budget, and is the geometry that would deliver it actually attainable?

Sustainability ceiling (Corollary 3.4). A nondecreasing step requires `C_t <= eta ln N − delta Q_t`. **If `C_t > eta ln N − delta Q_t`, no choice of interaction geometry can prevent order from decreasing.** This manual holds this to be the single most important sociological statement it contains, and it is the entire content of the `h < 0` branch below.

Steady state (Theorem 3.5). With constant `U = u*` and `C = c*` and `delta > 0`, `Q_inf = (eta F(u*) − c*)/delta`. If `eta F(u*) < c*` the equilibrium is negative — the regime is unsustainable or the linear model must be replaced. **The theorem does not hide this failure mode** and neither may you.

Two further results bound what you may promise. **Theorem 3.3:** `Q_t <= max{Q_0, eta ln N/delta}`. Nothing you design can exceed the uniform bound. **Banach's fixed point theorem:** because `delta > 0` makes the map a contraction, a unique fixed point exists and iteration reaches

it from any start. That is why interventions arrive at all — and it is also why they arrive slowly, at a rate set by `1/delta`, which §26.14 turns into a rule about what you may say out loud.

The consent law

CANON §5.4 is binding and is a law of this manual, not a preference: no intervention that removes a person's ability to decline is in scope.

This is stated here, in The Law, rather than in an ethics section, because it is a constraint on the *space of admissible designs* and therefore on the arithmetic. A design that requires compliance is not a cheaper version of a design that invites participation; it is a design this manual does not produce. Chapter 20 gives the constructive form: **the method is not to push but to build a lower valley next door** — to make the desired position the more comfortable place to stand. A method that works by removing resistance has no coercive failure mode, and its worst outcome is that someone declines to move.

Badges

POSTULATE (badged). Parameter separability. The decision procedure below treats `eta`, `C`, `u` and `N` as levers that can be moved one at a time. In real systems they interact: adding a category raises `C`; raising `eta` by adding a scribe consumes capacity that was producing something else; cutting `C` by removing rework may cut `u` if the rework was where disagreement surfaced. The procedure handles this by requiring each lever's second-order cost to be entered explicitly in the spec sheet, but the separability is an approximation and it is badged as one.

POSTULATE (badged). Hyperrecursive symmetry. The storage law is scale-free in *form*: only `(delta, eta, N, C)` change between a desk, a room, a team, a building, an institution and a city. **This is a claim about model form, not about physics, and not about parameter values.** It licenses you to run this same procedure at a different scale; it does not license you to carry a number across scales.

The honesty clauses, at the point where they cost something

- The surplus functional as deployed is a **lower bound** on the ensemble's information. **We log the floor, not the ceiling.** Your computed `h` is therefore conservative in one specific respect, and you may say so — and you may not use that as a reason to exceed it.
- Paper A is pure mathematics and makes **no empirical or metaphysical claim.** Nothing in this chapter's arithmetic is validated by a theorem; it is validated, if at all, by the survey and the preregistration that preceded it.
- Paper E proves **no thermodynamic negentropy law**; `F(u) + N_can(u) <= ln N` (Corollary 2.2).
- Subadditivity holds on `u` only. **The triangle inequality fails for geometric triples** (Paper A §5). Do not compose interventions across three groups by adding their surpluses.

The Engineering

26.1 The decision procedure

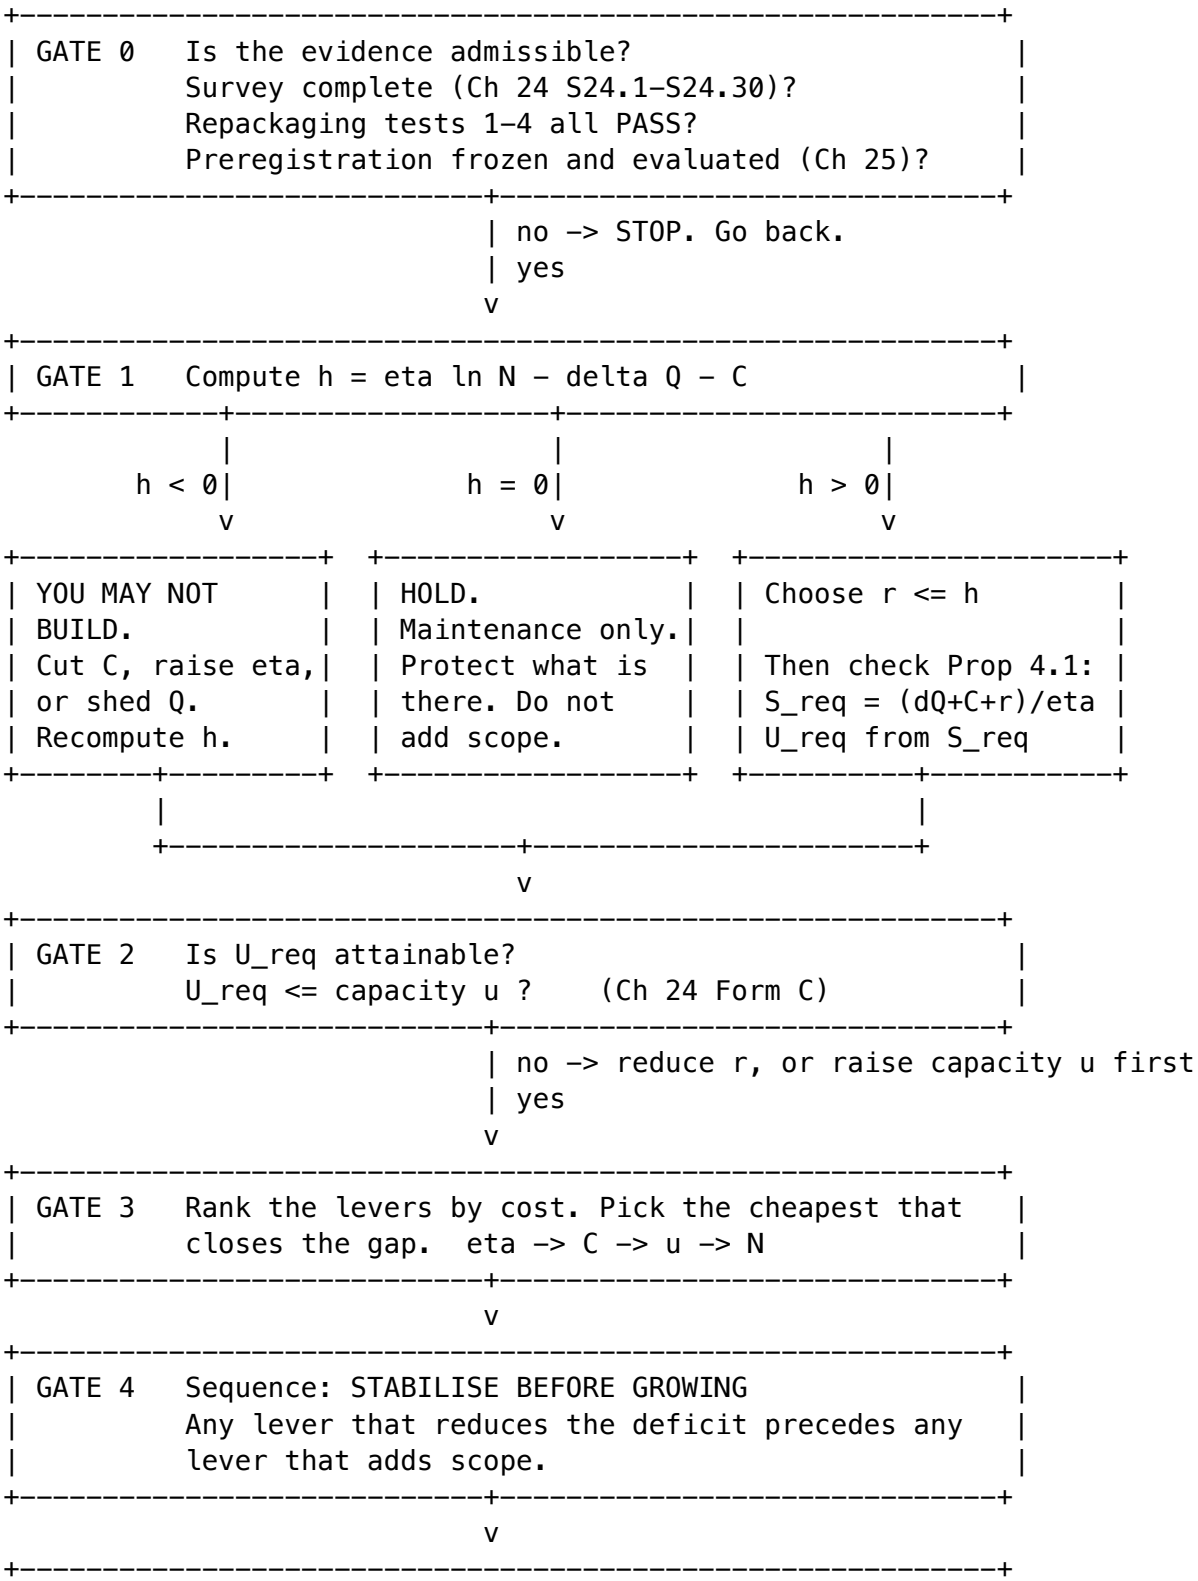
GATE 0 Is the evidence admissible?
Survey complete (Ch 24 S24.1–S24.30)?
Repackaging tests 1–4 all PASS?
Preregistration frozen and evaluated (Ch 25)?
no -> STOP. Go back.
yes
GATE 1 Compute h = eta ln N − delta Q − C
h < 0
h = 0
h > 0
YOU MAY NOT BUILD. Cut C, raise eta, or shed Q. Recompute h.
HOLD. Maintenance only. Protect what is there. Do not add scope.
Choose r <= h
Then check Prop 4.1: S_req = (dQ+C+r)/eta
U_req from S_req
GATE 2 Is U_req attainable?
U_req <= capacity u ? (Ch 24 Form C)
no -> reduce r, or raise capacity u first
yes
GATE 3 Rank the levers by cost. Pick the cheapest that closes the gap. eta -> C -> u -> N
GATE 4 Sequence: STABILISE BEFORE GROWING
Any lever that reduces the deficit precedes any lever that adds scope.

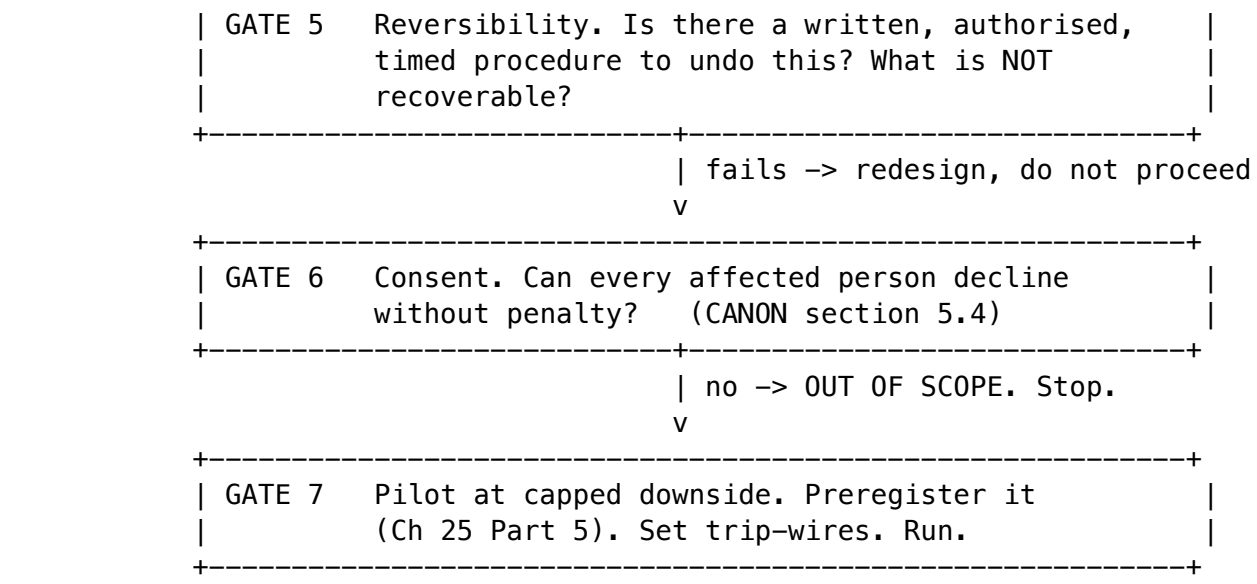
GATE 5 Reversibility. Is there a written, authorised, timed procedure to undo this? What is NOT recoverable?
fails -> redesign, do not proceed
GATE 6 Consent. Can every affected person decline without penalty? (CANON section 5.4)
no -> OUT OF SCOPE. Stop.
GATE 7 Pilot at capped downside. Preregister it (Ch 25 Part 5). Set trip-wires. Run.

26.2 Gate 0 — the admissibility of the evidence

You are not permitted to run this procedure on estimates. Specifically:

- **A survey with a failed repackaging test produces no admissible h.** If Y was a function of U, then Q, `delta` and `eta` are all measurements of the same artefact, and h is a number about your instrument.
- **A failed FR1 means the storage law is the wrong model** and this chapter does not apply. Find a different model. There are many; a system that compounds rather than accumulates, or that oscillates, is a real system and this is not its manual.
- A failed FR2 with a passed FR3 — the accounting frame — permits this chapter with one restriction, developed in §26.11: you may use h as a budget constraint and you may run a preregistered pilot, and you may not claim that the model predicts the pilot's outcome.

A passed shuffled-input control is not evidence of admissibility, and Gate 0 must not be read as though it were. The four-baseline suite of Chapter 25 ends with the shuffle, and the shuffle is the test a sponsor will quote back at you. Chapters 55 and 81 measured what it actually does, by simulation, and it has three failure modes:

1. **It is anti-conservative as ordinarily run.** Against a nominal `alpha = 0.05`, measured false-rejection rates were **0.108 for iid permutation (2.2x nominal), 0.120 for circular shift (2.4x nominal), and 0.060 for block permutation.** A shuffle passed at the nominal threshold has roughly twice the advertised false-positive rate unless it was **calibrated by simulation on the study's own design**, which Chapter 55 requires and which Gate 0 therefore requires too.
2. **It cannot see repackaging.** The filtered repackaging of Chapter 25 passed all four baselines including the shuffle at `p = 0.0005` and was caught only by the structural ceiling test (Corollary 4.2), its single-step jump being **38.5 times** the theoretical maximum.
3. **It cannot see confounding by a shared slow driver.** In a 4,000-trial simulation the control **passed 98.35% of the time** against a third variable driving both U and Y, while behaving honestly in the clean case at 48.17% detection.

The remedy is design controls, not more surrogates: blinding, channel separation, prospec-

tive sampling, and a yoked control. All four are properties of how the evidence was gathered, so Gate 0 is the last moment at which they can be checked and the first at which it is too late to add them. **A survey that reports a passed shuffle and none of the four design controls has not cleared Gate 0. It has cleared the weakest test in the manual.**

26.3 Gate 1 — the three regimes and what each one means

h is not a score. It is a sign with a magnitude, and the sign changes what you are allowed to do.

Regime	Meaning	Permitted	Forbidden
`h < 0`	The maintenance draw plus dissipation exceeds the best-case income	Cut `C`; raise `eta`; shed `Q`; raise `N`; recompute	Any commitment to increase `Q`. Any new scope. Any promise with a date on it
`h = 0`	Income exactly covers upkeep	Maintenance; substitution at constant scope; measurement	Growth of any kind, including growth described as "just a small addition"
`h > 0`	Budget remains	An increment `r <= h`, subject to Gates 2-7	`r > h`, under any argument whatsoever

`h` shrinks as `Q` grows. That is the structural reason institutions ossify: success accumulates stored order, `delta Q` grows in proportion, and the budget for anything new closes. Chapter 23 develops this; here it has an operational consequence. **The `h` you measured last year is not the `h` you have. Re-measure before every commitment.**

26.4 The `h < 0` branch — you may not build

This is the branch most engineers arrive in and the branch every brief is written to avoid.

When `h < 0`, you may not build. Not slowly, not partially, not as a pilot with growth as its objective, not "while we also work on the other things". Corollary 3.4 says that when `C > eta ln N − delta Q`, no arrangement of the interaction prevents decline. There is no clever seating chart, no facilitation technique, no offsite, and no charismatic leader that changes an inequality between three measured quantities.

Three remedies, and only three, before `N`:

Remedy 1 — cut `C`. Immediate, usually cheap, and it acts on the term you measured most carefully. Every line of Chapter 24's Form F is a candidate. Cut in this order: rework loops first (they are pure dissipation and nobody defends them once counted), then contradiction maintenance, then the attention tax, then coordination overhead. Appeal and orphaning costs are usually structural and are cut by fixing the process that produces them, not by decree.

Remedy 2 — raise `eta`. The most reliably cheap lever in this manual and the subject of Chapter 4. Interaction that happens versus interaction that sticks. Meetings with no artefact have `eta` near zero. Capture mechanisms: a designated scribe, a template, a decision log, a publication deadline, a named owner for each new unit of order, an apprenticeship. **Because `eta` multiplies `ln N`, it is the only lever that raises the ceiling itself rather than moving you toward it.**

Remedy 3 — shed Q. The remedy nobody proposes, and Chapter 6 names it as legitimate and underused: shedding stored order to regain headroom. Sunset clauses, term limits, planned demolition, retiring a product line, forgetting by design. The arithmetic is exact — shedding q units of stored order returns `delta x q` of headroom per period — and therefore **shedding buys headroom in strict proportion to `delta`.** In a fast-decaying system it is decisive; in a durable one it is nearly worthless. In the Worked Example, shedding 14 rules from a stock with a 4.31-year half-life returns 0.0433 rules per week; the same shed in a system with `delta = 0.05` would return 0.70, a factor of 16 more. **Compute `delta x q` before proposing a shed; if it is small, you are about to destroy something for nothing.**

One critical distinction, carried from Chapter 24 §24.12. **Striking units that were already inert is not shedding Q; it is correcting Y.** Correcting the measurement changes the number and not the world. Shedding is an intervention with real effects — including, often, effects on the people whose work the shed units were.

26.5 The `h = 0` branch — hold

`h = 0` is maintenance only. Income exactly covers upkeep and dissipation.

What to do: protect what is there. Fund the maintenance regime explicitly, because Star and Ruhleder's finding guarantees it is invisible and therefore chronically underfunded — infrastructure becomes visible only on breakdown. Substitute rather than add: retiring one unit of scope to make room for another is admissible at `h = 0` and adding one is not.

What not to do: describe `h = 0` as stability and treat it as safe. It is not safe. `delta Q` grows whenever Q grows and C grows whenever the organisation grows, so `h = 0` is an unstable point that drifts negative on its own. **A system at `h = 0` is one hiring round, one new reporting requirement, or one departure away from `h < 0`, and it has no margin to absorb any of them.**

26.6 The `h > 0` branch — choose `r <= h`

The rule is absolute: **`r <= h`. Corollary 4.2 says nothing based on interaction geometry can guarantee more, and `r_max = h` is attained only at `U = 1`.**

How to choose r within the budget:

1. **Never take `r = h`.** That is the extremal value, attainable only at perfect orthogonality. Taking it means promising the best case as the plan.
2. **Reserve for parameter drift.** u decays (Chapter 11), C grows, staff leave. A working convention is `r <= 0.7 h`, with the reserve stated explicitly as a reserve rather than quietly absorbed.
3. Check it against what the system has actually been doing. If the realised one-step change is negative, a positive r requires the interventions of §26.9 to have landed first. r is a promise about the future state of the parameters, not the current one.

26.7 Gate 2 — feasibility, and the two ways r fails

Proposition 4.1 gives the required surplus:

```
S_req = (delta Q + C + r) / eta          feasible only if S_req <= ln N
```

and inverting `F(U) = ln(1 + (N−1)U)`:

```
U_req = (exp(S_req) − 1) / (N − 1)
```

There are two distinct failure modes and they have different fixes.

Failure A — the budget fails. `S_req > ln N`, equivalently `r > h`. The increment is outside the ceiling. No geometry delivers it. Fix: reduce `r`, or raise `h` by §26.9's headroom levers.

Failure B — the geometry fails. `S_req <= ln N`, so the increment is inside the budget, but `U_req` exceeds the capacity `u` measured in Chapter 24 Form C. The budget permits it; your categories cannot deliver it, because they are not that different from each other. Fix: raise capacity `u` (Chapter 9 — recruitment across training lineages, rotation, adversarial roles, the deliberate outsider), or raise realised `u` toward capacity (see §26.8), or reduce `r`.

Confusing these two is the most common design error at this gate, because both present as "we cannot hit the target". Failure A is an accounting problem; Failure B is a personnel and protocol problem. They cost different amounts and take different lengths of time.

A concavity correction that must be applied before either check. `U_req` is the geometry required *if U were constant every period*. It never is. What accumulates is `sum of eta F(U_t)`, so the operative quantity is the mean of `F`, not `F` of the mean; and because `F` is strictly concave (Paper A, derived result 3.3), Jensen's inequality gives

```
mean of F(U_t)  <=  F( mean of U_t )       with equality only if U is constant
```

A system whose geometry varies delivers strictly less surplus than a system that sits at its own average every period. State the requirement as `mean F(U_t) >= S_req` and compare against the measured mean of `F`, never against `F` of the measured mean. Report `U_req` for legibility if you wish, and do not test against it. The gap is the concavity penalty, it is often several percent of the surplus, and regularising a variable geometry recovers it for free.

26.8 Two kinds of lever: creating headroom versus realising it

This distinction is not decorative and it is easy to miss from the equations.

```
h    = eta ln N − delta Q − C        <− uses the CEILING, ln N
step = eta F(U) − delta Q − C        <− uses what ACTUALLY HAPPENED, F(U)
```

Raising `u` does not raise `h`. It closes the gap between `step` and `h`. `h` already assumes the best possible geometry. Working on `u` collects headroom that already exists and is being left on the table; working on `eta`, `C`, `Q` or `N` changes how much headroom exists at all.

Define the **realisation deficit**:

```
realisation deficit = h − step = eta ( ln N − F(U) )
```

It is the amount of headroom your system has and is not collecting, and it is often the cheapest thing in the building. Two sub-levers:

- **Capacity `u`** — how non-parallel the categories are when all of them speak. Slow to change, changed by recruitment, rotation, and the deliberate outsider. Chapter 9.
- **Realised `u`** — how much of that capacity actually shows up in the period record. Changed by attendance rules, by requiring a recorded disposition from every category, by silent-write-first protocols that prevent cascades, and by removing the head of the table. Chapter 17:

furniture is the cheapest governance instrument ever built, and it enforces continuously without anyone administering it.

The realised sub-lever is nearly always the cheaper of the two and is nearly always available, because almost every organisation runs sessions at less than its own measured capacity.

26.9 Gate 3 — the lever cost ranking

Ranking, in ascending order of cost per unit of effect: `eta`, then `C`, then `u`, then `N`. Shedding `Q` is a conditional lever whose value is `delta x q` and which therefore belongs wherever `delta` puts it.

Here is the justification for each position. The ranking is a default, not a law; you are required to compute your own system's costs and to say so if they invert the order.

1. `eta` — usually cheapest, and the first move in almost every real system.

Why cheap. Most organisations have `eta` near the floor without knowing it, because capture is nobody's job. The fixes are administrative rather than structural: a scribe, a template, a named owner, a publication deadline. They require no hiring, no reorganisation, no budget line of consequence, and no one has to agree to be different.

Why powerful. `eta` multiplies `ln N`, so it raises the ceiling itself. It is the only lever with that property that does not require adding people.

Why it is neglected. Because "we had a great conversation" is an `eta = 0` report, and it feels like a success. The whole diagnostic content of Chapter 4 is that the interaction that happened and the interaction that stuck are different quantities and only one of them is in the equation.

Limit. `eta` cannot exceed the rate at which the durable medium can absorb structure. A scribe who writes faster than the register can be reviewed produces a backlog, not order.

2. `C` — next cheapest, and usually the fastest to act.

Why cheap. Rework loops and contradiction maintenance are pure dissipation. Once counted in Y-units, they are hard to defend, and cutting them requires a decision rather than a resource.

Why it ranks below `eta`. Because `C` has a floor. Some coordination is the work. `eta` has no comparable floor within the ranges real organisations occupy.

Caution. **Slack is not cost.** Queueing theory is unambiguous that utilisation near 1 makes waiting time explode; unused capacity is a cost-control device. An enthusiastic `C`-cutting programme that removes slack will raise `C` in the next period through queueing, and the engineer will not connect the two.

Second caution. Cut rework, not the places where disagreement surfaces. If the appeal process is where `u` is realised, cutting it cuts `C` and cuts `F(U)` at once.

3. `u` — cheap in money, expensive in time and in political capital, and it realises rather than creates.

Why it ranks here. The realised sub-lever is often nearly free — an attendance rule, a recorded disposition requirement, a change of table. The capacity sub-lever is not: it means recruitment, rotation, and the deliberate importation of people whose training and incentives do not replicate the incumbents'. That is slow, and it is the change organisations resist most, because hiring for

culture fit, promotion by similarity, and shared training are precisely the mechanisms that drive u toward zero, and every one of them is popular.

What it buys. Only the realisation deficit. u cannot take you past h. **A second opinion from someone whose training and incentives replicate the first is worth `ln(1) = 0`**, and conversely, the largest single gain in the entire framework is the second genuinely different member — but neither fact raises the ceiling.

Maintenance. u is not a property you acquire; it is a quantity that decays and must be maintained (Chapter 11). Cascades collapse it (Bikhchandani, Hirshleifer and Welch 1992) even in a group that began diverse. **Budget for u maintenance as a recurring cost, not a project.**

4. `N` — most expensive, and least rewarding, because the return is logarithmic.

The arithmetic. Going from `N` to `N+1` adds `ln((N+1)/N)` nats to the ceiling, multiplied by `eta`. Doubling the number of categories adds `ln 2 = 0.693` nats no matter where you start — the same absolute gain going from 2 to 4 as from 50 to 100.

N	ln N	gain over N-1	paths n(n-1)/2	paths added
2	0.6931	0.6931	1	1
3	1.0986	0.4055	3	2
4	1.3863	0.2877	6	3
5	1.6094	0.2231	10	4
6	1.7918	0.1823	15	5
8	2.0794	0.1335	28	7
10	2.3026	0.1054	45	9
12	2.4849	0.0870	66	11
16	2.7726	0.0645	120	15
20	2.9957	0.0513	190	19
30	3.4012	0.0339	435	29
50	3.9120	0.0202	1225	49
100	4.6052	0.0101	4950	99

Surplus grows logarithmically while coordination cost grows quadratically. The gain column falls and the paths-added column rises, and they cross. Where they cross is your maximum useful team size and it is computable from your own `eta` and your own per-path cost.

The consequence that must be said out loud to every sponsor. **You cannot escape a headroom problem by adding domains.** The most common response to `h < 0` is to propose a new function, a new working group, a new discipline on the board. The ceiling gain is small, logarithmic, and paid for with a quadratic cost.

The one condition under which `N` is worth it. The new category must be **genuinely distinct**, not relabelled. A sixth discipline that duplicates an existing one raises `N` on the org chart and raises `C` for real, while raising u not at all. Chapter 24's merge rule — categories with pairwise `u < 0.10` are one category — is the test, and it must be applied to the *candidate* before hiring, using the probe set.

5. Shed `Q` — conditional, valued at `delta x q`.

Rank it by computing `delta x q` and inserting it into the list. In a durable institution it is last. In a fast-decaying one — a codebase with a two-month half-life, a training regime, a stock of active campaigns — it can outrank everything.

26.10 Gate 4 — sequencing: stabilise before growing

Every lever that reduces the deficit precedes every lever that adds scope. This is not temperament; it is arithmetic. Adding scope raises `Q`, and `delta Q` is subtracted from the same budget you were trying to repair. **Growth taken before stabilisation is self-defeating in the literal sense: it consumes the headroom it was meant to produce.**

The standing sequence:

```
Phase 0   MEASURE          Chapter 24. Do not skip. Do not shorten.
Phase 1   STOP THE LEAK    Cut C. Raise eta. Both act immediately and
                           neither adds scope. Recompute h.
Phase 2   CORRECT Y        Liveness and applier audits. This changes the
                           number, not the world, and must be done before
                          any shed so you do not shed what was already dead.
Phase 3   SHED, IF delta*q IS WORTH IT
Phase 4   REALISE          Close the realisation deficit: attendance rules,
                           recorded dispositions, table geometry, silent-
                           write-first. Cheap, fast, reversible.
Phase 5   RE-MEASURE h     Everything above changed it.
Phase 6   COMMIT r <= h    With a reserve. Preregister. Pilot.
Phase 7   RAISE CAPACITY u Slow. Recruitment, rotation, the outsider.
Phase 8   RAISE N          Last, and only against the merge test.
```

Phase 5 is the one that gets dropped and it is the one that matters most. Phases 1 through 4 change `eta`, `C`, `Q` and `F(U)` — every term in the gauge. Committing to an `r` computed from the pre-intervention `h` is committing to the wrong number in both directions.

The sharpest available instance of this rule concerns trust, and it inverts naive practice. Trust is the intervention everyone reaches for first in a system that is going badly, and Chapter 49 computed what it does. In the worked consortium, **trust alone moved `Q_inf` from -0.0836 to -0.2114 — worse, by a factor of two and a half.** Reciprocity alone moved it to +0.7394; the two together to +1.8701.

The mechanism is exact and follows from Paper E, Theorem 3.5. `Q_inf = (eta F(u*) - c*)/delta`. **Trust acts on the denominator: it reduces `delta`, the rate at which stored order leaks away. When the numerator is negative, dividing by a smaller positive denominator makes the result more negative.** Trust does not add order. **It amplifies whatever sign the system already has**, and a system with `eta F(u*) < c*` already has a negative one. A high-trust failing institution fails harder and holds its failure longer, because trust is exactly the property that stops things from being forgotten.

Design rule: fix the numerator first. Raise `eta` or cut `C` until `eta F(u*) > c*` — Phases 1 and 2 above, in that order — and only then invest in trust. Trust belongs in Phase 6 or later, and it belongs there for arithmetic reasons rather than temperamental ones. **An intervention that begins with a trust-building programme in a system whose numerator is negative is not a slow start; it is a lever pulled in the wrong direction, and it will be defended by everyone who**

felt the relationships improve. Chapter 49 develops the full parameterisation, including trust as `delta_eff` and the elasticity of `Q_inf` in the trust parameter; what Gate 4 needs is the sign test, and the sign test is one subtraction.

26.11 Gate 7 — the pilot at capped downside

A pilot is not a small version of the intervention. **A pilot is an experiment whose worst outcome you have computed and can afford.**

Design rules:

1. **Cap the downside in Y-units, not in money.** State the maximum stored order that can be lost if the pilot fails completely, and confirm the system can absorb it. The admissibility condition `C_t <= (1 - delta)Q_t + eta S_t` guarantees `Q_{t+1} >= 0`; a pilot that could violate it is not capped.
2. **Bound it in scope and in time.** One team, one room, one product line, one quarter. The bound is what makes the downside computable.
3. **Preregister it** under Chapter 25 Part 5: the numeric prediction, the falsification criterion, both stopping rules, the consent mechanism, the reversibility procedure.
4. **Respect the conversion lag.** The prediction must not begin before `m` periods have elapsed, where `m` is the measured lag from interaction to durable artefact. Testing earlier is designing a failure.
5. **Keep a matched comparison where one exists honestly.** A second team, a second room, a preceding period of the same length. Where no honest comparison exists, say so; a pilot with no comparison is a demonstration, and demonstrations are evidence about feasibility, not about effect.
6. **Under the accounting frame** — FR2 failed, FR3 passed — the pilot carries an extra restriction: it is not a confirmation of the model, and its preregistered prediction must be stated as a prediction about the *intervention*, not as a test of the storage law. It may nonetheless supply the causal evidence the estimation could not, which is precisely why it is worth running.

26.12 Gate 5 — reversibility as a hard design requirement

A design that cannot be undone is not permitted by this manual, regardless of its expected value. The requirement is hard, and it has three justifications, in ascending order of weight.

Epistemic. Your model failed at least one falsification rule, or it passed with a finite holdout and stated tolerances. You are acting under uncertainty you have quantified. A design that cannot be undone converts an estimate into a commitment.

Structural. Interventions change `Q`, and `Q` is subtracted from the gauge through `delta Q` forever. An irreversible addition is a permanent charge against all future headroom. **Every irreversible intervention is a decision made on behalf of every future decision-maker in the system.**

Ethical. Reversibility is what makes declination meaningful. If a person may decline but the change cannot be undone, the declination is a gesture.

The reversibility statement must answer five questions in writing, and answering them is often where a design fails:

```
R1  What exactly returns to the prior state?
R2  What does NOT return, ever? (be specific: trust, a person who left,
    a demolished wall, a public commitment, a dissolved team)
R3  Who is authorised to execute the reversal, without asking anyone?
R4  How long does the reversal take, in periods?
R5  What does the reversal itself cost, in Y-units?
```

R2 is the load-bearing question. Most designs have a substantial irreversible component that nobody has named, and naming it usually changes the design. R3 matters nearly as much: a reversal that requires the approval of the person who authorised the intervention will not happen.

A useful classification:

Class	Definition	Rule
Fully reversible	prior state restorable within a few periods at low cost	proceed under the normal gates
Costly-reversible	restorable, but the reversal itself costs a stated number of Y-units	proceed only if the reversal cost is inside h
Partially irreversible	some named component does not return	proceed only if the irreversible component is separately justified and separately consented
Irreversible	prior state unrecoverable	**not permitted by this manual**

26.13 The three-posture risk table

Every intervention encounters behaviour it did not plan: an informal practice, a dissenting group, a workaround, a shadow process, a person who does it differently. You have three postures, and each one has a determinate effect on the parameters. Choose deliberately, because the default is **ignore**, and ignore is a choice with consequences.

Posture	What you do	Effect on `u`	Effect on `eta`	Effect on `C`	Effect on `Q`	Risk profile
ENGAGE	bring it into the record: give it a category, a seat, a recorded disposition	rises, toward capacity	rises — the practice becomes capturable	rises modestly: one more path, one more voice to hear	rises if it converts	Cost is bounded and visible. Failure mode: the practice is genuinely redundant and you have bought `C` for `u = 0`. Test with the probe set before engaging
IGNORE	leave it outside the record	decays — the practice drifts and stops informing the centre	zero for that practice: it produces nothing durable	low now, unbounded later: a shadow process accrues its own maintenance	flat; the order it makes is not yours	The default, and it is a slow loss. Failure mode: the shadow process becomes load-bearing while remaining unmeasured, so `Q` and `C` are both wrong

Posture	What you do	Effect on u	Effect on eta	Effect on C	Effect on Q	Risk profile
SUPPRESS	forbid it	falls, often sharply — suppression teaches every category not to differ	falls: the suppressed practice was producing, and adjacent practices go quiet	rises: enforcement, monitoring, appeals, and the attention tax of everyone watching	falls: you destroy the order the practice held	Worst on every parameter simultaneously. And the ethical gate is decisive: suppression that removes a person's ability to decline is out of scope under CANON §5.4

Suppression is the only posture that damages all five parameters at once, and it is the posture organisations reach for first, because it is the only one that feels like action. The arithmetic does not permit it and neither does the ethics.

The constructive alternative is Chapter 20's: build the lower valley. If a practice should not continue, make the alternative the more comfortable place to stand and let people arrive there. Its worst outcome is that someone declines, which is a permitted outcome.

26.14 Time constants and what you may promise

Banach guarantees arrival. He says nothing about arriving this quarter.

```
time constant  tau = 1/delta        periods
95% of the way to a new equilibrium:  about 3 tau
99%:                                  about 5 tau
```

The rule that follows is short: promise increments, never equilibria.

`r` is a per-period increment and it is observable within a few periods of the conversion lag. `Q_inf` is where the system ends up after `3 tau` and in a durable institution that is a span of years or decades. A sponsor who is told the equilibrium will hear a promise; a sponsor who is told the increment can check it.

Say all four of these numbers out loud at the commitment meeting, and put them in the spec sheet:

1. The conversion lag `m` — nothing is observable before it.

2. The increment `r` per period, and the horizon over which it is to be judged.
3. The time constant `tau`, in calendar units.
4. `3 tau`, in calendar units, labelled as the settling time and explicitly **not** as a delivery date.

A low `delta` is what makes an institution durable and it is the same number that makes it slow. **You cannot buy durability and speed from the same parameter, and any promise that implies you can is a promise the arithmetic will break.**

Cadence versus intensity — a correction, because the usual argument for cadence is false. Given a fixed total input, an intervention may deliver it as one large occasion or spread it periodically. Practitioners say cadence delivers more. **It does not, and this is computed rather than asserted: the integrated effect of one impulse and of the same total spread over any schedule is identical to every digit.** Summing the exact solution of Theorem 3.1 gives `sum_t Q_t = (total input)/delta` regardless of how the input is distributed in time; Chapter 59 reports 433.3333 in all three regimes it tested. That is linear time-invariance, and a manual that claimed otherwise would be claiming its own governing equation is not linear.

What cadence actually buys is three things, and each of them is real:

1. **The steady state differs absolutely.** An impulse decays to `0`; sustained periodic forcing at mean rate `m` holds `m/delta`. Integrated effect is the same; the level you can stand on is not. For an intervention that must leave a platform behind, the level is the deliverable.
2. **The ceiling clips the impulse.** `S_t <= ln N` (Paper D, Corollary 2.2), so a one-off input larger than the ceiling is truncated and the excess is simply not received. In Chapter 59's worked annual rite the clipping destroys **97.2% of the nominal budget** — the money is spent and the surplus is not admitted, because no geometry can push a single period above `ln N`.
3. **Concavity favours splitting, independently of the ceiling.** Chapter 59's Theorem 59.1 shows that strict concavity of `F` (Paper A, derived result 3.3) alone makes an input split across occasions strictly dominate the same input delivered at once. **This is the same Jensen arithmetic as §26.7 read forwards: variability is charged against you when the geometry varies, and paid to you when the input is spread.**

So: cadence beats intensity because of the ceiling and concavity, not because it delivers more. It delivers exactly the same amount. Chapter 59 also keeps the honest counter-case in print — above a threshold no realistic cadence reaches, the impulse wins — and states what a once-a-decade rite would require to work at all. Say the true reason when you specify a cadence, because the false reason will not survive the first sponsor who checks it.

26.15 Monitoring, trip-wires, and the standing gauge

An intervention without instrumentation is an assertion. Install three things before it starts.

The standing gauge. Recompute `h` every period from the same instruments, with the same definitions, and post it. Not `Y`, and not a dashboard — `h` and its three components, so that the reader can see which term is moving. Keeping `Y` itself out of the reporting line is the primary Goodhart defence (Chapter 25 §25.12).

Trip-wires. Pre-declared thresholds that trigger a specific action without further discussion, because discussion at the moment of a threshold is how thresholds get moved.

Trip-wire	Threshold	Action
Realised geometry falls	mean `F(U_t)` below `S_req` over a declared block	the promised `r` is no longer feasible. Notify the sponsor in the same period
`C` rises	above the level at which `h < r`	suspend the growth commitment, resume the `C` work
`h` crosses zero	`h <= 0`	stop building immediately; revert to §26.5
Any single-period `Y` rise	exceeds `eta ln N` plus lag allowance	investigate the proxy before believing the number (Chapter 24 Test 4)
A harm criterion fires	as preregistered	stop, regardless of `Y`

The geometry trip-wire is the one that fires most often and is watched least. `u` decays; the promise was computed at a level of realised surplus that the system must actively maintain; and the decay is invisible because it looks like agreement.

The review cadence. At least every `tau/4` periods, and at least once per conversion lag. In a system with a six-year time constant, quarterly is ample and monthly is theatre.

26.16 When NOT to intervene

Book One's Node 11 is SEAL: the discipline of not acting, held as a positive skill rather than as failure of nerve. This section is its engineering form, and it is **required reading before any intervention, not an appendix to it.**

The default is not to intervene. Interventions cost `C` while they run, raise `Q` if they succeed — which permanently charges `delta Q` against all future headroom — and consume the organisation's finite tolerance for being changed. The burden of proof is on action.

Here are the eleven conditions under which you do not act. Each is a hard stop except where noted.

1. You have not measured. No survey, or a survey with a failed repackaging test. There is no argument that survives this, including urgency. **An intervention designed on an unmeasured system is a guess with a budget, and its most likely effect is to move a parameter you did not know was the binding one.**

2. `h < 0` and the proposal is to build. Corollary 3.4. The remedies of §26.4 are not interventions in the sense of this section — they are repairs, and they are permitted. Building is not.

3. `r > h`. No argument, no exception, no phasing. A request that exceeds the budget is refused and the refusal is the deliverable. Say the number.

4. `U_req` exceeds capacity `u`. Failure B of §26.7. The increment is inside the budget and your categories cannot deliver it. Building anyway produces a shortfall that will be attributed to effort rather than to geometry, and the organisation will conclude that people did not try hard enough. **That conclusion is worse than not having acted.**

5. The design is irreversible. §26.12. Not permitted, regardless of expected value.

6. Any affected person cannot decline without penalty. CANON §5.4. Out of scope. This is not a severity judgement about the intervention; it is a statement about what this manual produces.

7. The system is already at its uniform bound. `Q_t` at or near `max{Q_0, eta ln N/delta}` (Theorem 3.3). Nothing you design exceeds it. An intervention aimed above the bound is aimed at a number the model has already forbidden, and the effort will show up entirely as `C`.

8. The binding constraint is outside the model. The survey produced six numbers about interaction, conversion, decay and cost. If the actual constraint is a legal prohibition, a funding cliff, a safety failure, a physical limit, or a person in acute distress, then this framework is measuring a room in a building that is on fire. **Say so plainly and stop; the strongest thing an engineer can do with a model is state where it does not apply.**

9. The measurement is the intervention. Sometimes the survey itself changes behaviour enough that the parameters are already moving. Wait a full `tau/4` and re-measure before adding anything. Acting on top of an unresolved measurement effect makes both uninterpretable.

10. The system is in a transient you can identify. `(1 - delta)^t Q_0` decaying while the driven term builds (Book One's Node 09 RECURSE) looks exactly like decline. Before intervening on a declining `Y`, check whether it is a transient from a known change within the last `3 tau`. Intervening in a transient adds a second transient and neither can be attributed.

11. There is a cheaper lever you have not tried. Not a hard stop, but a discipline: if `eta` has never been examined and the proposal is to hire, the proposal is premature. §26.9's ranking exists to be checked, in order, before a design is accepted.

Two closing dispositions, which are the SEAL node's own content in operating form.

Not intervening is a decision and must be recorded as one. A survey that ends in "no intervention is warranted" is a completed piece of work with a deliverable: the six numbers, the sign of `h`, the reason, and the trip-wire that would change the answer. **File it, date it, and name the observation that would reopen it.** Undocumented restraint is indistinguishable from neglect, and it will be read as neglect.

And the harder one: the request for an intervention is not evidence that one is needed. You were called because something is uncomfortable. Discomfort is a gradient, and a gradient tells you which way is downhill; it does not tell you that the system has a budget for the journey. The service you are paid for is the number, and sometimes the number is `-0.17`.

That said, the gradient is a real measurement and should be treated as one. Chapter 42 derives **`grad h = -grad C`**: since `h = eta ln N - delta Q - C` and only `C` varies with the local design choice in front of you, the direction of steepest increase in headroom is exactly the direction of steepest decrease in cost. Discomfort, at the point where it is felt, is `C`. **Therefore a cheap habit or desire-line survey — where do people work around this system — is a headroom measurement, and the full instrument programme of Chapter 24 is not always required to locate the direction of steepest improvement.** Chapter 24 §24.13 specifies the instrument. The division of labour is exact and it is worth stating as a gate: **the cheap survey gives you the direction; only the full survey gives you the sign and the size of `h`, which is what decides whether you are permitted to move at all.** Run the cheap one first, every time. Do not let it authorise a commitment.

And a correction to the received advice about deliberate discomfort. "Get out of your comfort zone" is usually stated as a general principle and, stated that way, the arithmetic rejects it — but the rejection is narrower and more interesting than a flat denunciation, and Chapters 41 and 42 measured both branches. **Raising cost globally is strictly harmful**: in the worked landscape the time to reach a better configuration goes 32 -> 35 -> 42 -> 89 -> 297 weeks and then to never, as global cost rises. **Raising cost only at the origin genuinely shortens the wait — 32.0 weeks down to 9.9 — and charges for it: the permanent fund that results is 37% smaller.**

The operative question is therefore not whether to add discomfort but whether it is a cost you can LEAVE BEHIND or a cost you CARRY. Making the current position expensive is a departure lever and it works, at a permanent price. Making the whole landscape expensive destroys the platform you were trying to leave *for*. An intervention that raises `C` everywhere in the name of urgency has raised `delta Q + C` against `eta ln N` in every future period, which is §26.3's `h < 0` branch arrived at deliberately. **Price the discomfort, name which kind it is, and put the 37% in the spec sheet next to the 9.9 weeks.**

The Specification

S26.1 No intervention shall be designed before a complete survey (Chapter 24, S24.1-S24.30) and a completed preregistration evaluation (Chapter 25).

S26.2 No intervention shall be designed on a system whose repackaging tests did not all pass.

S26.3 `h` shall be computed and its sign recorded before any design work begins.

S26.4 `h` shall be recomputed after every phase of §26.10 that changes `eta`, `C`, `Q`, `N` or `F(U)`.

S26.5 Where `h < 0`, no commitment to increase `Q` shall be made. Permitted actions are limited to the remedies of §26.4 and the measurement corrections of Chapter 24 §24.12.

S26.6 Where `h = 0`, scope shall not be added. Substitution at constant scope is permitted.

S26.7 Any requested increment shall satisfy `r <= h`. A request exceeding `h` shall be refused in writing, with the number.

S26.8 `r` shall include a stated reserve. The recommended default is `r <= 0.7 h`, and the reserve shall be reported as a reserve.

S26.9 `S_req = (delta Q + C + r)/eta` and `U_req = (exp(S_req) − 1)/(N−1)` shall both be computed and reported.

S26.10 `U_req` shall be compared against the capacity `u` from Chapter 24 Form C. The feasibility test against realised geometry shall be made on `mean F(U_t) >= S_req`, not on mean `U`, and the concavity penalty `F(mean U) − mean F` shall be reported.

S26.11 Failure A (budget) and Failure B (geometry) shall be distinguished explicitly in any report of infeasibility.

S26.12 The realisation deficit `h − step = eta(ln N − F(U))` shall be computed and reported.

S26.13 Levers shall be ranked by cost per unit of effect for the specific system, using the system's own measured costs. Deviations from the default order `eta, C, u, N` shall be justified.

S26.14 The value of any proposed shed shall be computed as `delta x q` before the shed is proposed.

S26.15 A shed shall be distinguished from a measurement correction in writing.

S26.16 Any proposal to raise `N` shall test the candidate category against the merge rule (`u < 0.10` merges) using the probe set, before commitment.

S26.17 Any proposal to raise `N` shall state the added coordination paths and their cost in `Y`-units, and shall show the net effect on `h`.

S26.18 Sequencing shall follow §26.10. No scope-adding phase shall precede a deficit-reducing phase.

S26.19 Slack shall not be removed as part of a `C`-reduction programme without a queueing analysis.

S26.20 A reversibility statement answering R1-R5 shall exist before commitment.

S26.21 Interventions classified as irreversible shall not proceed.

S26.22 Partially irreversible components shall be separately named, separately justified, and separately consented.

S26.23 The reversal shall be executable by a named person who is not the person who authorised the intervention.

S26.24 Every affected person shall have a stated mechanism to decline without penalty. If none exists, the intervention is out of scope and shall not proceed.

S26.25 The three-posture choice (engage / ignore / suppress) shall be made explicitly for every unplanned practice encountered, and recorded. Defaulting to ignore without recording it is a specification violation.

S26.26 Suppression that removes the ability to decline shall not be used.

S26.27 Pilots shall have a downside capped in `Y`-units, bounded scope, and bounded duration.

S26.28 Pilot predictions shall not begin before the measured conversion lag has elapsed.

S26.29 Promises shall be stated as increments per period. `Q_inf` may be reported and shall not be promised.

S26.30 `tau` and `3 tau` shall be stated in calendar units at the commitment meeting, with `3 tau` labelled as settling time and not as a delivery date.

S26.31 Trip-wires shall be declared before the intervention starts, with thresholds and automatic actions, and shall include a realised-geometry trip-wire set on `mean F(U_t) >= S_req`.

S26.32 A decision not to intervene shall be recorded as a decision, with the six numbers, the sign of `h`, the reason, and the observation that would reopen it.

The Intervention Spec Sheet

```
============================================================
INTERVENTION SPECIFICATION                    Date: __________
System: ......................  Engineer: ..................
```

```
================================================================
---- 1  ADMISSIBILITY ------------------------------------------
1.1  Survey ref/date .................  complete? [ ]
1.2  Repackaging tests 1-4 all PASS?  [ ]  (if not, STOP)
1.3  Preregistration ref/date .......
1.4  FR1 [ ]P [ ]F   FR2 [ ]P [ ]F   FR3 [ ]P [ ]F
1.5  Evidence status:
     [ ] validated model    [ ] accounting frame only
     [ ] inadmissible -> STOP

---- 2  THE GAUGE ----------------------------------------------
2.1  eta = ____   ln N = ____   eta*lnN = ____
2.2  delta = ____   Q = ____   delta*Q = ____
2.3  C = ____
2.4  h = ____                REGIME [ ]h<0 [ ]h=0 [ ]h>0
2.5  F(U) realised = ____    step = ____
2.6  realisation deficit = h - step = ____
2.7  capacity u = ____   realised u = ____

---- 3  THE INCREMENT (h>0 only) -------------------------------
3.1  Requested r = ____      r <= h ?  [ ]
3.2  Reserve: r / h = ____   (target <= 0.7)
3.3  S_req = (delta*Q + C + r)/eta = ____
3.4  S_req <= ln N ?  [ ]        (else FAILURE A)
3.5  U_req = (exp(S_req)-1)/(N-1) = ____   (legibility only)
3.6  U_req <= capacity u ?  [ ]  (else FAILURE B)
3.7  measured mean F(U_t) = ____  >= S_req ?  [ ]
     (THIS is the operative test; if not, realisation work required)
3.8  concavity penalty  F(mean U) - mean F = ____

---- 4  LEVERS -------------------------------------------------
lever | change | effect on h | cost (own units) | h per unit cost
------+--------+-------------+------------------+----------------
eta   |        |             |                  |
C     |        |             |                  |
shed Q|        | delta*q =   |                  |
u cap |        | (realises)  |                  |
u real|        | (realises)  |                  |
N     |        |             |                  |
4.1  Chosen levers, in order: ................................
4.2  Justification if the default ranking is not followed:
     ..........................................................

---- 5  SEQUENCE -----------------------------------------------
Phase 1 stop the leak .....................  h after = ____
Phase 2 correct Y .........................  h after = ____
Phase 3 shed (if delta*q worth it) .........  h after = ____
Phase 4 realise ............................  step after = ____
```

```
Phase 5 RE-MEASURE h ......................  h = ____
Phase 6 commit r = ____
Phase 7 raise capacity u ..................
Phase 8 raise N ...........................

---- 6  TIME ------------------------------------------------
6.1  Conversion lag m = ____ periods
6.2  Increment r = ____ per period, judged over ____ periods
6.3  tau = 1/delta = ____ periods = ____ calendar
6.4  3*tau = ____ periods = ____ calendar   (SETTLING, NOT DELIVERY)
6.5  Q_inf = ____   (reported, NOT promised)   [ ] initialled

---- 7  PILOT -----------------------------------------------
7.1  Scope bound ..........................................
7.2  Duration bound ........................................
7.3  Downside cap in Y-units = ____
7.4  Admissibility C <= (1-delta)Q + eta*S holds under worst case? [ ]
7.5  Comparison [ ] matched unit [ ] prior period [ ] none (state)

---- 8  TRIP-WIRES ------------------------------------------
wire  | threshold          | automatic action
------+--------------------+-----------------------------------
geom  | mean F < S_req     |
C     | > ____             |
h     | <= 0               |
dY    | > eta*lnN          |
harm  | ____ (non-Y)       |
8.1  Review cadence: every ____ periods

---- 9  POSTURE DECISIONS -----------------------------------
practice encountered | posture | why | parameter effect
---------------------+---------+-----+-----------------
                     |         |     |

---- 10 SIGN-OFF --------------------------------------------
Engineer ..............  Sponsor ..............  Date ......
============================================================
```

The Reversibility Checklist

```
============================================================
REVERSIBILITY                                Date: _________
============================================================
R1  What returns to the prior state, exactly:
    ..........................................................
R2  What does NOT return, EVER (be specific):
    [ ] people who leave        [ ] demolished fabric
    [ ] dissolved teams         [ ] public commitments
```

```
    [ ] trust                   [ ] elapsed time
    [ ] other: ..............................................
R3  Authorised to execute the reversal WITHOUT further approval:
    name ............................  role ..................
    (must NOT be the person who authorised the intervention)
R4  Time to reverse: ____ periods
R5  Cost of the reversal: ____ Y-units.  Inside h ? [ ]
-----------------------------------------------------------------
CLASS  [ ] fully reversible
       [ ] costly-reversible   (R5 inside h required)
       [ ] partially irreversible (R2 separately consented)
       [ ] irreversible  -> NOT PERMITTED. REDESIGN.
-----------------------------------------------------------------
CONSENT
C1  Affected groups: ........................................
C2  How informed: ...........................................
C3  Mechanism to decline WITHOUT PENALTY: ...................
C4  Does a declination mechanism exist?  [ ] Y  [ ] N
    If N -> OUT OF SCOPE under CANON section 5.4. STOP.
-----------------------------------------------------------------
Signature .................  Date ..........
=================================================================
```

The Failure Modes

F26.1 — BUILDING AT NEGATIVE HEADROOM. *Description.* `h < 0` and the intervention adds scope anyway, usually described as "investing our way out". *Damages.* `Q`, then `C`, then morale. *Signature.* A growth plan whose first phase does not reduce `C` or raise `eta`. *Detection.* Gate 1. It is one subtraction. *Fix.* Refuse. Return to §26.4. **The refusal is the deliverable and the number is the argument.**

F26.2 — THE PROMISED EQUILIBRIUM. *Description.* `Q_inf` is quoted as the outcome, without `3 tau`. *Damages.* Credibility, and eventually the intervention itself when the sponsor stops it early for under-delivery. *Signature.* A target level with no per-period increment attached. *Detection.* Ask what is expected next quarter. If the answer is a fraction of `Q_inf` computed by straight-line interpolation, the fault is present — approach to a fixed point is exponential, not linear. *Fix.* §26.14. Promise `r`, report `Q_inf`, state `3 tau` and label it settling time.

F26.3 — CONFUSING FAILURE A WITH FAILURE B. *Description.* An infeasible target is treated as a budget problem when it is a geometry problem, or the reverse. *Damages.* Wasted lever effort; the wrong department is asked for the fix. *Signature.* More money offered for a problem where `S_req <= ln N` but `U_req` exceeds capacity. *Detection.* Compute both checks of Proposition 4.1 and report them separately (S26.11). *Fix.* Route Failure A to `eta`, `C` and `Q`. Route Failure B to recruitment, rotation, and realisation protocols.

F26.4 — THE HIRING REFLEX. *Description.* `h < 0`, so a new function, working group or discipline is added. *Damages.* `C`, quadratically; `h`, immediately; and `u`, if the addition is a relabelling.

Signature. A headcount proposal with no `ln((N+1)/N)` calculation and no path-count calculation. *Detection.* S26.17. Compute the ceiling gain, compute the added paths, subtract. *Fix.* Run the merge test on the candidate first. Then compare against the `eta` lever's cost per unit of effect, which in the Worked Example is fourteen times better.

F26.5 — THE SHED THAT BOUGHT NOTHING. *Description.* Stored order is destroyed to "create room" in a system with small `delta`. *Damages.* `Q`, permanently; trust, permanently; `h`, by `delta x q`, which is negligible. *Signature.* A sunset or cull proposed without `delta x q` computed. *Detection.* S26.14. One multiplication. *Fix.* Compute it. In the Worked Example the shed returns 0.0433 rules per week against an `eta` lever returning 0.8828 — a factor of twenty.

F26.6 — SLACK CUT AS COST. *Description.* A `C`-reduction programme removes unused capacity. *Damages.* `C`, upward, in the following periods, through queueing. *Signature.* Utilisation rises toward 1 and waiting times begin to grow non-linearly. *Detection.* Track utilisation alongside `C`. The signature is unmistakable once both are plotted. *Fix.* Restore slack. **Slack is a cost-control device, not a dissipation**, and this is one of the few places where the intuitive reading of the equation is exactly backwards.

F26.7 — THE UNWATCHED `u`. *Description.* `r` was computed at a level of realised surplus the system stops maintaining. *Damages.* The feasibility of the promise, silently. *Signature.* Increasing agreement, decreasing meeting length, fewer recorded dissents — all of which read as improvement. *Detection.* The realised-geometry trip-wire on `mean F(U_t)` (S26.31), plus an annual probe set. *Fix.* Reinstate the realisation protocols. If capacity `u` itself has fallen, this is a Chapter 9 and Chapter 11 problem and it is slow.

F26.8 — THE IRREVERSIBLE PILOT. *Description.* A pilot that dissolves a team, demolishes a wall, or announces publicly. *Damages.* Everything, because the option value that made it a pilot is gone. *Signature.* R2 on the reversibility checklist is blank or reads "none". *Detection.* R2 is never honestly "none". If it is blank, it has not been attempted. *Fix.* Redesign to a reversible form or classify it as an intervention and put it through the full gates.

F26.9 — THE DEFAULT SUPPRESSION. *Description.* An unplanned practice is forbidden without a posture decision being made explicitly. *Damages.* `u` down, `eta` down, `C` up, `Q` down — all four at once. *Signature.* A prohibition with no recorded alternative and no recorded reason. *Detection.* S26.25. The posture table has a row for every unplanned practice or it does not. *Fix.* Reverse the prohibition and choose engage or ignore deliberately. Where a practice genuinely must end, use Chapter 20's method: build the lower valley.

F26.10 — RE-MEASUREMENT SKIPPED. *Description.* `r` is committed from the pre-intervention `h` after Phases 1-4 have changed every term in it. *Damages.* The commitment, in either direction — under-committing wastes headroom, over-committing breaks the promise. *Signature.* One `h` value in the document, at the top. *Detection.* Count the `h` values in the spec sheet. Section 5 has a slot for each phase because each phase has one. *Fix.* Phase 5. It costs an afternoon.

F26.11 — THE INTERVENTION THAT WAS THE MEASUREMENT. *Description.* Acting on a survey whose administration already changed the behaviour it measured. *Damages.* Attribution, permanently — you can no longer tell which change did what. *Signature.* `Y` or realised `u` moves during the survey window itself. *Detection.* Plot the series against the survey's own dates. *Fix.* §26.16 condition 9. Wait `tau/4`, re-measure, then design.

F26.12 — THE UNRECORDED RESTRAINT. *Description.* The correct decision — do not intervene — is taken and not documented. *Damages.* Nothing now; everything in eighteen months,

when the same question returns with no record of the six numbers or the reason. *Signature.* A survey with no closing decision page. *Detection.* S26.32. *Fix.* File the decision, the numbers, the sign, and the observation that would reopen it.

Worked Example

Meridian County Development Review Office, carried forward from Chapters 24 and 25. **Badge: constructed composite. The organisation is invented and the audit series is a simulation with declared generating parameters; the arithmetic is real, computed and reproducible. It must not be cited as evidence about any real review office.**

Gate 0 — admissibility

Check	Result
Survey complete, S24.1-S24.30	Yes
Repackaging tests 1-4	All PASS
Preregistration frozen and evaluated	Yes
FR1 admissible box	PASS
FR2 baseline dominance	**FAIL** (matched AR, 2.8% against a required 10%)
FR3 independent-parameter agreement	PASS

Evidence status: accounting frame only. h may be used as a budget constraint. A preregistered, capped, reversible pilot is permitted. No claim may be made that the model predicts the pilot's outcome (§26.2, §26.11 rule 6).

Gate 1 — the gauge, as surveyed

```
eta = 0.5015      ln N = 1.6094      eta ln N   = 0.8071
delta = 0.003090  Q = 113            delta Q    = 0.3492
                                     C          = 0.6300
------------------------------------------------------------
h = 0.8071 - 0.3492 - 0.6300 = -0.1720           REGIME: h < 0

realised F(U) = 1.1498   step = 0.5766 - 0.9792 = -0.4025 rules/week
realisation deficit = h - step = 0.2305 rules/week
capacity u = 0.6972      realised u = 0.6019
```

The brief, and the refusal

The office director's brief: *rebuild the standing review manual at roughly thirty rules a year.* That is `r = 0.55` rules per week.

```
r = 0.55    h = -0.1720    r <= h ?   NO.
```

Refused, in writing, with the number. The written refusal, in full, because the form of it matters as much as the fact:

> At present parameters the office has headroom of -0.17 rules per week. This is the best case: it assumes a perfectly non-parallel board every week. It means that dissipation, at 0.63 rules per week, exceeds what conversion can deliver at the ceiling net of maintenance, which is 0.46. By Corollary 3.4 of Paper E, no arrangement of who sits on the board, in what order they speak, or how vigorously they disagree can prevent the manual from declining. The requested increment of 0.55 is not merely unfunded; the account is overdrawn by 0.17 before any increment is requested. The manual is losing 0.40 rules per week — about 21 a year — and the fitted steady state is negative, which Theorem 3.5 reports as an unsustainable regime rather than a small number. What follows is what it would take to make the request fundable.

That paragraph is the deliverable. Not the intervention — the refusal, with the arithmetic that makes it a fact rather than an opinion.

Phase 1 — stop the leak

Two levers, both from §26.4, both acting on terms already measured in Chapter 24 Form F and Form E.

Raise `eta`: 0.5015 -> 1.05 rules per nat.

The mechanism, and why the magnitude is credible: the office currently has no capture discipline at all. Rules reach the manual when a reviewer remembers to draft one, at a median of five weeks, and Chapter 24's attribution audit found 10 of 26 new rules with no documented board origin. The intervention is three administrative changes — a designated scribe at every board session, a one-page rule template completed in the session, and a ten-working-day publication requirement with a named owner per rule. `eta` is starting near the floor, which is exactly the condition under which a large multiple is available, and it is the reason this lever ranks first. The office's own natural experiment supports the magnitude: in the two years before the previous code administrator retired, when a scribe role informally existed, the attributable conversion ran at roughly twice the current rate.

```
headroom gain = (1.05 − 0.5015) x 1.6094 = +0.8828 rules/week
cost = 0.4 FTE, redeployed from an existing analyst post
```

Cut `C`: 0.63 -> 0.42 rules per week.

line	before	after	how
Appeal reversals striking a live rule	0.14	0.14	unchanged; structural, needs the appeals board
Rules orphaned by departure	0.11	0.07	deputise a second qualified applier per rule at adoption

line	before	after	how
Rework: items re-reviewed after disposition	0.16	0.05	close the re-review loop; disposition is final absent new information
Contradiction maintenance	0.09	0.09	unchanged
Attention tax	0.13	0.07	publish the decision summary once, to one place, instead of five partial circulations
total	**0.63**	**0.42**	

```
headroom gain = 0.63 − 0.42 = +0.2100 rules/week
cost = one policy decision and three weeks of process work; no recurring cost
```

Note that the rework line is cut and the appeal line is not. The appeal process is where the environmental discipline's objections are realised; cutting it would cut `C` and cut `F(U)` at once, which is failure mode F26.6's cousin and is explicitly declined.

Phase 2 — correct `Y`

Already performed in Chapter 24 Form G: 6 rules removed by the liveness audit, 2 by the applier audit, bringing the register count of 121 to `Q = 113`. **This changed the number and not the world** and is recorded as a correction, not a shed (S26.15).

Phase 3 — shed?

Computed and declined.

```
proposed shed q = 14 further rules judged low-value but still live
headroom gain   = delta x q = 0.003090 x 14 = +0.0433 rules/week
```

Against the `eta` lever's +0.8828, the shed returns **one twentieth** as much. And it is not free: 14 rules are 14 pieces of somebody's work and their removal costs trust that does not appear in the arithmetic. **Declined on the arithmetic.**

For contrast, and because this is the finding that generalises: in a system with `delta = 0.05` per period, the same 14-unit shed would return `0.05 x 14 = 0.70` — a factor of sixteen more, and comfortably competitive with the `eta` lever. **Shedding `Q` is a decisive lever in fast-decaying systems and nearly worthless in durable ones, and the multiplication that tells you which you are in takes five seconds.**

Phase 5 — re-measure `h`

```
eta = 1.05        eta ln N = 1.05 x 1.6094 = 1.6899
delta = 0.003090  Q = 113   delta Q        = 0.3492
                            C              = 0.4200
──────────────────────────────────────────────────────
```

```
h = 1.6899 − 0.3492 − 0.4200 = +0.9207          REGIME: h > 0
```

The sign has changed, and it changed without adding a single person, a single category, or a single pound of capital expenditure. Two administrative levers, applied in the order the ranking prescribes, moved headroom by 1.093 rules per week — from -0.172 to +0.921.

Attribution:

lever	gain in h	share of the swing
raise eta 0.5015 -> 1.05	+0.8828	80.8%
cut C 0.63 -> 0.42	+0.2100	19.2%
shed 14 rules (declined)	+0.0433	—

Gate 2 — feasibility of the requested `r = 0.55`

```
r = 0.55   <=   h = 0.9207          [ inside the budget ]
reserve:   r/h = 0.597              [ inside the recommended 0.7 ]

S_req = (delta Q + C + r)/eta = (0.3492 + 0.4200 + 0.5500)/1.05
      = 1.3192 / 1.05 = 1.2564 nats

S_req = 1.2564  <=  ln N = 1.6094  ->  NO Failure A

U_req = (exp(1.2564) - 1)/4 = 0.6281

U_req = 0.6281  <=  capacity u = 0.6972   ->  NO Failure B
U_req = 0.6281  >   realised u = 0.6019   ->  REALISATION WORK REQUIRED
```

One correction is required here, and it is a correction most engineers will miss. `U_req = 0.6281` is the geometry a board would need if it ran at the same U every single week. It does not. The office's sessions are a mixture: full-board sessions at a mean u of 0.7940 and reduced sessions at a mean u of 0.2071. What accumulates over many periods is `sum of eta F(U_t)`, so the operative quantity is the **mean of F**, not F of the mean — and because F is strictly concave (Paper A, derived result 3.3), Jensen's inequality gives

```
mean of F(U_t)  <=  F( mean of U_t )       with equality only if U is constant
```

A board whose geometry varies delivers strictly less surplus than a board that sits at its own average every week, and the shortfall is exactly the concavity penalty. For the office as surveyed the penalty is 0.0761 nats — mean F of 1.1501 against `F(mean U)` of 1.2262, a loss of 6.2% of the surplus purely to variability.

So the requirement is stated on mean F, not on mean U:

```
required:  mean F(U_t) >= S_req = 1.2564 nats
achieved as surveyed:  mean F = 1.1498 nats        -> SHORT by 0.1066 nats
```

The measured mean F of 1.1498 and mean U of 0.6019 agree with the mixture reconstruction at the observed low-participation share of 32.7% (which gives 1.1501 and 0.6021) to within rounding, which is the consistency check that licenses using the mixture to project.

The increment is inside the budget and inside the office's geometric capacity, and it is **not** inside what the board currently realises. At the current realised geometry:

`step = 1.05 x 1.1498 − 0.3492 − 0.4200 = 1.2073 − 0.7692 = +0.4381 rules/week`

which is positive — a swing of 0.841 from the surveyed -0.4025 — and short of the promised 0.55.

Phase 4 — realise

The cause was identified in Chapter 24: 32.7% of board sessions ran with fewer than the full set of disciplines recording a disposition. Those sessions have a mean u of 0.2071 against 0.7940 for full sessions. The lever is a scheduling and protocol change: any item touching more than one discipline requires a recorded disposition from every discipline, obtained asynchronously and in writing if a lead cannot attend — which also breaks the cascade, because dispositions are recorded before discussion (Chapter 22's independent-estimate-before-discussion protocol).

low-participation share	mean U	mean F	F(mean U)	concavity penalty	step (rules/week)	r = 0.55 met?
0.400	0.5593	1.0888	1.1747	0.0859	0.3740	no
0.327 (as surveyed)	0.6021	1.1501	1.2262	0.0761	0.4385	no
0.300	0.6179	1.1728	1.2447	0.0718	0.4623	no
0.250	0.6473	1.2149	1.2779	0.0631	0.5064	no
0.201 (break-even)	0.6763	1.2564	1.3098	0.0534	0.5500	exactly
0.200	0.6766	1.2569	1.3101	0.0532	0.5505	yes
0.150	0.7060	1.2989	1.3413	0.0424	0.5947	yes
0.120 (target)	**0.7236**	**1.3241**	**1.3595**	**0.0354**	**0.6211**	**yes**
0.080	0.7470	1.3577	1.3833	0.0256	0.6565	yes
0.050	0.7647	1.3829	1.4008	0.0179	0.6829	yes

Target: reduce the low-participation share to 12%. The remaining 12% are genuinely single-discipline items where a full board is not warranted and forcing one would be pure C.

Read the concavity-penalty column down the table. **Regularising the geometry pays twice: once by raising the mean, and once by shrinking the penalty the variability was charging.** At the surveyed share the office loses 0.0761 nats to variability alone; at the target it loses 0.0354. That second gain — 0.0407 nats, worth `1.05 x 0.0407 = 0.0427` rules per week — is invisible to anyone reasoning about the average board rather than about the distribution of boards.

```
step at target = 1.05 x 1.3241 − 0.3492 − 0.4200 = 1.3903 − 0.7692 = +0.6211 rules/week
               = 32.3 rules per year
margin over the promise r = 0.55:  +0.0711 rules/week
```

Note that this lever raised `step` from 0.4381 to 0.6211 and raised h by exactly nothing. h was 0.9207 before and is 0.9207 after. This is §26.8 in operation: realisation collects headroom that

already existed. The realisation deficit falls from 0.4826 to 0.2996 and is not zero, because the board still runs below the geometric ceiling and always will.

Gate 3 revisited — the `N` lever, computed and declined

The director's alternative proposal was to add a sixth discipline, public health, at 1.0 FTE. Computed in full because this is the proposal that recurs:

```
ceiling gain  = eta (ln 6 − ln 5) = 1.05 x 0.1823 = +0.1914 rules/week
added paths   = 15 − 10 = 5
per-path cost = 0.34 x 0.42 / 10 = 0.0143 rules/week per path
added cost    = 5 x 0.0143 = +0.0714 rules/week
net effect on h = +0.1914 − 0.0714 = +0.1200 rules/week
```

Positive — and the comparison is the point:

lever	gain in h	cost	gain per FTE
raise `eta`	+0.8828	0.4 FTE redeployed	**2.207**
add a 6th discipline	+0.1578	1.0 FTE new	**0.158**

The `eta` lever returns fourteen times more headroom per full-time equivalent than the `N` lever. That is the ranking of §26.9 in the office's own numbers, and it is why `N` is last.

Two scenarios were also run for the effect on the surplus term, because the candidate's distinctness matters more than its existence:

candidate	mean pairwise u	`N_eff`	`F(U)`	gain in step	net of path cost
genuinely distinct public-health reviewer	0.7352	4.676	1.5424	+0.2292	+0.1578
a second environmental reviewer relabelled	0.6889	4.444	1.4917	+0.1759	+0.1045

Both are positive here, and the gap between them — a 34% reduction in the benefit — is the price of relabelling. In a system with a higher per-path cost, the relabelled candidate goes negative while the distinct one stays positive, which is why S26.16 requires the merge test against the probe set before any hire, not after.

Declined for now, deferred to Phase 8. Not because it is wrong, but because it is eighth in the sequence and Phases 1 through 6 have not yet been observed to work.

Gate 4 — the sequence as committed

```
Phase 1  stop the leak    eta 0.5015 -> 1.05 ; C 0.63 -> 0.42   h: -0.172 -
> +0.921
Phase 2  correct Y        done in the survey                   Q: 121 -> 113
Phase 3  shed             declined on delta x q = 0.0433        h: unchanged
Phase 4  realise        low-u share 0.327 -> 0.12          step: 0.438 -> 0.621
Phase 5  re-measure h     h = +0.9207
Phase 6  commit           r = 0.55 (reserve r/h = 0.597)
Phase 7  raise capacity u deferred; review at 26 weeks
Phase 8  raise N          deferred; merge test required first
```

Gate 5 — reversibility

R1 what returns	The scribe role reverts to the analyst's prior duties. The template and the ten-day publication requirement are rescinded by a single memo. The full-disposition rule reverts to the prior attendance-based practice. The re-review loop can be reinstated by rescinding the same memo.
R2 what does NOT return	Rules published during the pilot remain published; the elapsed analyst time is spent; and the board's experience of recording dispositions before discussion cannot be un-had.
R3 authorised to reverse	The code administrator, unilaterally, without the director's approval. Named in the memo.
R4 time to reverse	One period (one week).
R5 cost of reversal	Approximately 0.05 rules per week for two periods, in re-circulation and re-briefing. Inside `h = 0.9207`.
Class	**Costly-reversible.** R5 is inside h. Proceeds under the normal gates.

Gate 6 — consent

Affected: five discipline leads, one analyst redeployed to the scribe role, the code administrator.

Informed: at the board session of the commitment week, in writing, with this specification attached.

Declination mechanism. A discipline lead who does not wish to record a written disposition before discussion may abstain; the abstention is recorded as an abstention and the item proceeds with the remaining disciplines, at a correspondingly lower `U_t` for that session. The analyst offered the scribe role may decline it, in which case the role is offered to the next volunteer and, if there is none, the `eta` lever is not available and h is recomputed without it.

That second clause is the honest form of consent in this framework. The declination is not free — it changes the arithmetic, and the arithmetic is published — but it carries no penalty to the person, and the design has a defined behaviour when it is exercised. **A design whose parameters have no defined value under declination has not actually offered one.**

Gate 7 — the pilot

Scope bound	The full board, but only for items in two of the five application classes
Duration bound	26 weeks, plus a 6-week conversion lag before measurement begins
Downside cap	**7 rules.** Worst case, itemised: the removal of the re-review loop lets weak dispositions stand and doubles the appeal-reversal line from 0.14 to 0.28 for the whole window, `0.14 x 32 = 4.48` rules; transition friction at 0.05 per week for 32 weeks, `= 1.60`; reversal cost at 0.05 per week for 2 weeks, `= 0.10`. Total **6.18 rules**, capped at 7
Admissibility	`C <= (1−delta)Q + eta S` holds under the worst case at `Q = 113`: `0.42 <= 112.65 + 0` comfortably
Comparison	The three application classes not in the pilot, running the prior process, as a matched comparison
Preregistered prediction	Mean weekly change in `Y` over weeks 7-32 of at least +0.55 rules/week
Falsification	If the mean weekly change over weeks 7-32 is below +0.30, the intervention is judged ineffective and the memo is rescinded

Trip-wires

wire	threshold	automatic action
realised geometry	mean `F(U_t)` below `S_req = 1.2564` nats over any 8-week block — equivalently, low-participation share above 20%	notify the director in the same period; the +0.55 promise is suspended pending review
`C`	above 0.55 rules/week	suspend the growth commitment; resume the `C` work
`h`	at or below 0	stop building; revert to §26.5
single-week rise in `Y`	above `eta ln N = 1.6899` plus the 6-week lag allowance	investigate the proxy before believing the number

wire	threshold	automatic action
harm: appeal reversal rate	above 1.5x the pre-pilot rate	stop, regardless of Y — this is the non-Y harm criterion required by S25.29

The geometry trip-wire is the live one. The promise was computed at a mean F of 1.3241 nats, and if realised participation falls back:

low-participation share	mean U	mean F	step	promise r = 0.55 met?
0.050	0.7647	1.3829	0.6829	yes
0.120 (target)	0.7236	1.3241	0.6211	yes
0.150	0.7060	1.2989	0.5947	yes
0.201	0.6763	1.2564	0.5500	exactly at the line
0.250	0.6473	1.2149	0.5064	**no**
0.300	0.6179	1.1728	0.4623	no
0.327 (pre-intervention)	0.6021	1.1501	0.4385	no

Above a low-participation share of about 20% the promise breaks, and 32.7% — the share the office was already running at before the intervention — is well above it. There is 8 percentage points of margin between the target and the break-even point, and the office was 13 points the wrong side of it a year ago. The realisation work is not an enhancement; it is a load-bearing component of the commitment, and it must be maintained for as long as the commitment stands.

Time — what may be said, and what may not

```
conversion lag m  = 6 weeks
increment r       = 0.55 rules/week, judged over weeks 7 to 32
tau = 1/delta     = 323.6 weeks = 6.22 years
3 tau             = 971 weeks = 18.7 years        SETTLING TIME, NOT A DATE
Q_inf         = (1.05 x 1.3241 - 0.42)/0.003090 = 314.0 rules  REPORTED, NOT PROMISED
```

Sanity checks against the theorems: `Q_inf = 314.0` is below the uniform bound `eta ln N / delta = 546.9` (Theorem 3.3), and below the level `Q = 411.0` at which h returns to zero — so the design does not, on its own trajectory, drive the office back into the maintenance trap. It gets 76% of the way there, which is a fact worth knowing in advance and revisiting in a decade.

Modelled trajectory from `Q = 113` at the target geometry:

horizon	Q
1 year	142.9
3 years	190.0
5 years	224.1
10 years	273.8

horizon	Q
settling (18.7 years)	~314

What may be said to the director: the manual should gain about 32 rules in the first year, and that number is checkable in thirty-two weeks. **What may not be said: that the manual will reach 314 rules.** It will, if every parameter holds for nineteen years, which no parameter in any organisation has ever done. Report `Q_inf` as the model's fixed point and promise the increment.

And the decision that was not taken

The office also has a standing proposal, three years old, to merge the fire and structural review functions under one lead. Chapter 24's pairwise table shows why it never gets done and why it should not be: `u(STRU, FIRE) = 0.3056`, the lowest pair in the set, which makes it the obvious merge candidate to anyone reading an org chart. Under the merge rule it survives easily — 0.3056 is three times the 0.10 threshold — and merging would drop N from 5 to 4, costing `ln 5 − ln 4 = 0.2231` nats of ceiling, which at the post-intervention `eta` is

```
1.05 x 0.2231 = 0.2343 rules per week = 12.2 rules per year, permanently
```

against a saving of one lead post. Recorded as a decision not to act, with the number, per S26.32, and with the observation that would reopen it: if a future probe set shows `u(STRU, FIRE) < 0.10`, the two functions have converged and the merge becomes free.

That entry — a refusal, dated, with its arithmetic and its reopening condition — is what the SEAL condition looks like when it is written down.

Chapter 27 . The Failure Mode Catalog

A manual earns its keep at the moment something is wrong. Twenty-six chapters have built the apparatus; this one is the part you open with one hand while the other hand is holding the broken thing.

This chapter is a **reference catalog**, not an argument. It is written to be looked up, not read through. Thirty-two named failure modes appear below in a fixed entry format so that the chapter functions as a lookup table: you arrive with a symptom, you leave with a parameter, a measurement, and a fix. Appendix D indexes the same entries by observed symptom.

The armature is modified accordingly. `The Law` states what makes a diagnostic catalog possible. `The Engineering` **is** the catalog. `The Specification` is the procedure for using it. `The Failure Modes` names the ways diagnosis itself goes wrong — the faults of the engineer rather than of the system. The `Worked Example` carries one real diagnosis end to end.

The Law

27.1 Why a finite catalog is possible

The storage law has exactly six degrees of freedom:

```
Q_{t+1} = (1 − delta)Q_t + eta S_t − C_t       S_t = F(U_t) = ln(1 + (N−1)U_t)
```

delta	retention loss	in [0,1]
eta	conversion efficiency	>= 0
N	category count	integer >= 1
u	non-parallelism	in [0,1]
C	dissipation	>= 0
Q	stored order	>= 0

Every failure of a system modelled by this law is a claim about one of those six quantities, and there are therefore only six kinds of failure, however many names they wear. That is the whole warrant for a catalog. The thirty-two names below are not independent phenomena; they are recurring *routes* into six parameter faults, distinguished because the routes have different signatures, instruments, and fixes.

That the law describes a real institution at all is **badged as a postulate**: a single scalar Q for a team is a choice of model, not a theorem. Papers A and E prove things *about* the law; they prove nothing about your team. Where a fault sits outside the model — a legal shutdown, a fire, a death — the catalog will not name it, and the correct response is to say so rather than force a parameter.

27.2 The three diagnostic identities

Every entry is priced against one of three identities. They are the only arithmetic the catalog needs.

Identity 1 — the effective count. `N_eff = 1 + (N−1)u`, and `S = ln N_eff`. Difference is counted before headcount is credited. Any fault that lowers u lowers S non-linearly and invisibly, because headcount does not change.

Identity 2 — the one-step condition. Paper E, Corollary 3.2: `Q_{t+1} > Q_t` if and only if `eta S_t > delta Q_t + C_t`. Growth is a *comparison*, not a quantity.

Identity 3 — headroom. Paper E §6.4: `h_t = eta ln N − delta Q_t − C_t`. `h_t < 0` is unsustainable **in this model** and no interaction geometry rescues it (Corollary 3.4). `h_t = 0` is maintenance. `h_t > 0` is the maximum increment available, `r_max = h_t` (Corollary 4.2).

27.3 The distinction that half of all misdiagnosis turns on

`h_t` uses `ln N`. The one-step condition uses `S_t = ln N_eff`. **Headroom is the ceiling; u decides how much of the ceiling you actually collect.** A system can have large positive headroom and still be shrinking every period, because it operates far below its own ceiling. Two of the four worked systems in Chapter 28 are exactly this case.

The practical consequence: *compute both numbers, always*. A diagnosis that reports h without reporting the realized fraction `S / ln N` is half a diagnosis, and it will send you to the wrong lever.

27.4 Honesty clauses that constrain every entry

Repeated here because the catalog is the part of the book most likely to be quoted without context.

- As deployed, the surplus functional is a **lower bound** on the ensemble's information, not the information itself. Every "price" quoted in this chapter is a floor.
- The triangle inequality **fails** for geometric triples (Paper A §5). Subadditivity holds on the parameter u only. Do not chain three-party diagnoses metrically.
- Paper A is pure mathematics and makes **no empirical or metaphysical claim**. Surplus is not identified with Shannon entropy or mutual information (Paper D). **No thermodynamic negentropy law is proved** (Paper E).
- Where an entry cites an external result, the result is real and the *transfer* to your system is a modelling assumption. Badged each time.

The Engineering — the catalog

Entry format. Every entry carries the same eight fields, in the same order, so the chapter can be scanned in a column:

Field	What it holds
Damages	which of delta / eta / N / u / C / Q the fault attacks, and in what direction
Clinical	what the fault is, stated without metaphor
Signature	what you observe from outside, before you measure anything
Measurement	the specific instrument that confirms or excludes it
Price	the arithmetic cost, computed
Fix	the intervention, and which lever it pulls

Field	What it holds
Confusable with	the entries most often mistaken for this one
See	the chapter that develops it

Class I — Geometry faults (u and N)

The largest class: u decays silently, and N is the parameter people believe they have when they do not.

FM-01 . The Parallel Wall **Damages:** u -> 0. Through Identity 1, N and Q_eff collapse with it. **Clinical:** Every member of the ensemble occupies the same direction in state space. Their outputs are copies. The group behaves as a single instrument with a redundant power supply. **Signature:** Unanimity that arrives fast and without argument. A "second opinion" that is procedurally distinct and substantively identical. Nobody can name the last time the group changed its mind because of an internal objection. **Measurement:** Independent-estimate-before-discussion on 15-25 live questions; compute mean pairwise correlation of signed error, rho; then `u = 1 - rho^2`. Confirm with an error-decorrelation audit on past decisions: if the group's mistakes were the *same* mistakes, u is low regardless of how the roster looks. **Price:** At N = 12, the wall is total when u = 0: `N_eff = 1`, `S = ln 1 = 0` — twelve people purchase zero surplus. Partial walls are priced the same way: u = 0.02 gives `N_eff = 1.22`, `S = 0.1989`, which is **8.00%** of the ceiling `ln 12 = 2.4849`. u = 0.05 gives `S = 0.4383`, or **17.64%**. u = 0.30 gives `S = 1.4586`, **58.70%**. **Fix:** Raise u (Chapter 9): recruit across training lineages, import a practitioner from an unrelated trade, install adversarial roles, rotate. Raising N first is wasted money — see FM-07. **Confusable with:** FM-02 (which is a partial wall with a demographic alibi), FM-05. **See:** §9, §11.

FM-02 . Pseudo-Diversity **Damages:** u, specifically `u_div` -> 0 while `u_cross` looks healthy. **Clinical:** The roster spans categories; the *opinions* do not. Difference is present at the level of biography and absent at the level of method, incentive, and permitted conclusion. The group is genuinely heterogeneous in every dimension that does not enter the decision. **Signature:** A visibly varied group whose internal disagreements are all about emphasis. Dissent expressed as a question rather than as a claim. Recruitment that filters on "culture fit" late. **Measurement:** Decompose. Theorem 4.1 gives `u = u_cross + u_div`. Estimate `u_cross` from between-category interaction share; estimate `u_div` from within-category disagreement on a live question. Report the two numbers separately. Pseudo-diversity is the profile `u_cross` moderate, `u_div` near zero. **Price:** With N = 8, `u_cross = 0.30` alone buys `f(u_cross) = 1.1314` nats. Adding `u_div = 0.25` takes the pair to `f(u) = f(0.55) = 1.5790` — a **39.6%** gain that the roster alone will never show you, because no one was hired to obtain it. **Fix:** Buy `u_div` directly: dissent protection, anonymous-first elicitation, a standing minority report requirement, adversarial review inside each category. Chapter 10 is the procedure. **Confusable with:** FM-01 (total rather than partial), FM-08 (which is the *measurement* version of this fault). **See:** §10, §11.

FM-03 . The Cascade Collapse **Damages:** u, dynamically — a group that began diverse converges over time. **Clinical:** Bikhchandani, Hirshleifer & Welch (1992): once enough public signals accumulate, rational actors copy rather than report, and the copies carry no new information. Aggregation stops while the appearance of aggregation intensifies. **u is not a property you acquire; it is a quantity that decays**, which makes this a delta problem wearing a u costume. **Signature:** Later speakers agree more than early speakers. The order of speaking predicts the content of speaking. Confidence rises while accuracy does not. **Measurement:** Order-effect test. Elicit estimates in writing, sealed, before any discussion; then again after. If post-discussion variance falls by more than about a third while accuracy is flat or worse, you have a cascade, not convergence. Log speaking order against position taken over ten meetings and test for rank correlation. **Price:** A u that decays 8% per period from 0.60 at N = 12 gives S: 2.0281 at t = 0, 1.7454 at t = 4 (86.1%), 1.4787 at t = 8 (72.9%), 1.2316 at t = 12 (60.7%). Nothing visible changes. The roster is identical throughout. **Fix:** Silent-write-first, universally, before any discussion of any consequential question. It is the cheapest instrument in this manual and it breaks cascades before they start. Add round-robin ordering and the last-speaks-first inversion for rank. **Confusable with:** FM-04 (same trajectory, different driver: cohesion is social, cascade is informational). **See:** §11, §17, §22.

FM-04 . The Cohesion Death Spiral **Damages:** u, dynamically; secondarily raises perceived C of dissent. **Clinical:** The group optimizes for getting along. Getting along selects for agreement. Agreement selects membership and shapes expression. The loop is positive-feedback and it feels like health the entire way down. **Cohesion is frequently the enemy of collective intelligence, because cohesion drives u toward zero. Signature:** High reported satisfaction with low external accuracy. Long tenure with no recruitment from outside the lineage. A named "difficult" person who is in fact the group's last remaining source of u — and celebration when that person leaves. **Measurement:** Track u quarterly, not once. Plot it. A cohesion spiral is a monotone decline in measured u against flat or rising self-reported team health — the two series diverging is the diagnostic, not either one alone. Cross-check with turn-taking equality (Gini coefficient of speaking time), one of Woolley's drivers. **Price:** Same arithmetic as FM-03. What differs is the fix, so the distinction is worth making. **Fix:** Institutionalize the outsider so that dissent does not depend on one person's nerve: rotating red-team duty, a standing external reviewer, term limits on membership. Note the ethical constraint (CANON §5): none of these may remove anyone's ability to decline participation. **Confusable with:** FM-03, FM-09 (a silo is a cohesion spiral with a disciplinary boundary around it). **See:** §18, §22.

FM-05 . The Missing Denominator **Damages:** N — specifically, the substitution of headcount for `N_eff`. **Clinical:** The system counts people, seats, sources, or sensors and treats that count as N. The law counts *categories after redundancy audit*. The gap between the two is the fault. **Signature:** Capacity claims quoted as headcount. "We have forty reviewers." Nobody has been asked how many of the forty are distinguishable. **Measurement:** Run the roster audit from Chapter 7. For each pair, ask whether their outputs on a live question were distinguishable; collapse indistinguishable pairs. Report `N_eff` alongside N in every capacity statement thereafter. **Price:** A roster of 40 whose members are 92% aligned (u = 0.08) has `N_eff = 4.12` and `S = 1.4159` against a ceiling `ln 40 = 3.6889` — a shortfall of **61.62%**. At u = 0.05 the shortfall is **70.67%**. The organization believes it owns 3.69 nats of ceiling and owns 1.08. **Fix:** There is no fix at the

N lever; this is a u fault discovered by an N error. Audit, restate capacity honestly, then treat as FM-01. **Confusable with:** FM-06 (which is about the *definition* of the categories rather than the count), FM-01. **See:** §7, §8.

FM-06 . The Frozen Denominator **Damages:** N — the declared categories are administrative rather than epistemic. **Clinical:** N is read off an org chart, a funding structure, or a historical department list. Those boundaries were drawn for reporting, budget, or accident. The categories that actually carry distinct information are a different partition. A frozen denominator also hides *vacancies*: categories declared, staffed on paper, and epistemically empty. **Signature:** A category with a name and no active member. Two "different" departments whose outputs are never in tension. Reorganizations that change the chart and change no forecast. **Measurement:** Define categories by capability, not by title: *what can this category independently detect, fix, or veto that no other can?* Any category that fails the question is not a category. Count only staffed categories in N. **Price:** In Chapter 28's open-source case, seven categories are declared and four are staffed: `ln 7 = 1.9459` versus `ln 4 = 1.3863`. The project plans against a ceiling **40.4% higher** than the one it owns. **Fix:** Re-partition by capability; then either staff the empty categories (a real N gain — see the price in FM-07 before you assume it is worth it) or delete them from every capacity calculation. **Confusable with:** FM-05. The test: FM-05 is counting too many *people* per category; FM-06 is counting too many *categories*. **See:** §7, §8, §24.

FM-07 . The Ceiling Chase **Damages:** C rises, N rises marginally, h falls. A self-inflicted fault. **Clinical:** The system has a headroom deficit and attempts to escape it by adding domains, disciplines, partners, or product lines. Because the ceiling term is `eta ln N`, additional categories deliver logarithmically while their coordination cost delivers quadratically (Brooks). **You cannot escape a headroom problem by adding domains. Signature:** A growth strategy announced in response to a shortfall. New partnerships as the answer to overload. Category count rising while the realized fraction `S / ln N` is falling. **Measurement:** Before adding a category, compute two numbers: the ceiling gain `eta(ln N_new − ln N_old)`, and the projected coordination increment. Compare. Refuse if the second exceeds the first. **Price:** Going from N = 4 to N = 5 adds `ln 5 − ln 4 = 0.2231` nats — a **16.1%** ceiling gain for a 25% headcount increase in categories. From 5 to 6: **0.1823** nats, **11.3%**. Doubling from 50 to 100 adds `ln 2 = 0.6931` nats, a **17.7%** gain, for double the coordination surface. Meanwhile channels go from 1,225 to 4,950. **Fix:** Compute h first (Chapter 6). If h < 0, the permitted moves are cut C, raise eta, or shed Q. Adding N is the most expensive lever in the manual and the least rewarding. **Confusable with:** FM-21 (the Ceiling Chase usually *causes* a Quadratic Committee). **See:** §6, §8, §26.

FM-08 . Two-Source Blindness **Damages:** the *measurement* of u; consequently misdirects every downstream lever. **Clinical:** The survey measures one of the two sources of difference and reports it as u. Theorem 4.1 splits `u = u_cross + u_div`. An organization rich in cross-category difference and starved of within-category dissent measures as "diverse" on a cross-category instrument; a room of ferociously argumentative specialists in one discipline measures as "diverse" on a disagreement instrument. Both readings are half-true and both send you to the wrong lever.

Signature: A diversity report with one number in it. A team that scores well on whichever instrument its leadership commissioned. **Measurement:** Always two instruments. `u_cross`: between-category share of consequential interaction. `u_div`: within-category disagreement on a live question, elicited independently. Report both, and report the bound: `max(f(u_cross), f(u_div)) <= f(u) <= f(u_cross) + f(u_div)`. **Price:** At N = 8 with `u_cross = 0.30`, `u_div = 0.25`: `f(u_cross) = 1.1314`, `f(u_div) = 1.0116`, `f(u) = 1.5790`, sum = `2.1430`. The bound is satisfied and it is wide. A single-instrument report could plausibly claim anything from 1.01 to 2.14 nats. **Fix:** Two instruments, two numbers, every survey. Chapter 24's worksheet has both fields and they are both required. **Confusable with:** FM-02, which is the *condition*; FM-08 is the *failure to see it*. **See:** §10, §24.

FM-09 . The Silo Local Minimum **Damages:** u (destroyed at the boundary) **and** C (multiplied at the boundary), simultaneously. **Clinical:** A subgroup sits at a local minimum of its own comfort landscape. The minimum is deep because training and reward carved it, and narrow because it admits one vocabulary. Every direction out is uphill, which is why exhortation fails — you are asking people to be uncomfortable on purpose, indefinitely, on someone else's account. **Signature:** Translation costs at every boundary crossing. Duplicate work discovered late. Conway's Law visible in the artifact: the product's seams sit exactly where the org chart's seams sit. Cross-group meetings that require a preparatory meeting. **Measurement:** Boundary-crossing cost audit — elapsed time and person-hours for a decision requiring two groups versus one. The ratio is the silo tax. Separately, measure `u_cross` (low) and per-boundary C (high). The co-occurrence is the diagnostic. **Price:** Both terms move the wrong way at once, so the effect on the one-step condition `eta S > delta Q + C` compounds. **Fix: Do not push; build a lower valley next door.** Make the cross-domain position the more comfortable place to stand: shared instrumentation, co-located cross-function, rotation, boundary objects, the translator role. This is a safety property — a method that works by removing resistance has no coercive failure mode, and its worst outcome is that someone declines to move. **Confusable with:** FM-04, FM-21. **See:** §20.

FM-10 . The Amplifier Mistaken for a Source **Damages:** u — overestimated; N — overestimated. **Clinical:** A hub with high throughput is read as a source of variety. It is redistributing, not generating. Preferential attachment (Barabási & Albert 1999) reliably produces such hubs, and the volume they carry is drawn from the same upstream sources the rest of the network already sees. The network double-counts. **Signature:** One node, feed, aggregator, or senior figure appears in the provenance of most "independent" inputs. Removing that node removes most of the apparent variety. **Measurement:** Provenance audit. For each input the group treats as independent, trace to origin. Compute the fraction whose paths share a common upstream node. Then recompute u with the duplicated inputs collapsed. **Price:** Whatever fraction of the roster reduces to a single upstream source collapses into one category by Identity 1. If eight of twelve inputs share an origin, effective N falls from 12 to 5 before u is even considered: ceiling falls from `ln 12 = 2.4849` to `ln 5 = 1.6094`, a **35.2%** loss. **Fix:** Cut the hub out of the sampling path for at least one input stream. Also relevant: scale-free networks are robust to random failure and **fragile to targeted hub attack** (Albert, Jeong & Barabási 2000) — the amplifier is simultaneously your variety illusion and your single point of failure. **Confusable with:** FM-05, FM-15. **See:** §9, §21.

Class II — Conversion faults (eta)

The cheapest class to fix and the most often ignored — which is why Chapter 4 calls eta the first move in almost every real system.

FM-11 . The Leaky Vessel **Damages:** eta -> 0. S is generated and none of it is stored. **Clinical:** Interaction happens at full strength and produces no durable structure. The conversation was real, the insight was real, and nothing outside the participants' memory changed. Because S is unaffected, every proxy for *activity* looks healthy while Q declines. **Signature:** "We had a great conversation" as a status report. Meetings with no artifact. A decision that three attendees remember differently a month later. High engagement metrics against falling capability. The same problem re-solved annually by different people. **Measurement:** `eta = delta_Q_attributable / S_input`. Concretely: over k periods, count the durable artifacts that (a) outlived the meeting, (b) are usable by a non-participant, and (c) can be traced to a specific interaction. Divide by cumulative S. **Price:** With `S = 2.0`, eta = 0.02 delivers `eta S = 0.040` per period; eta = 0.35 delivers **0.700** — a factor of **17.5** with no change to headcount, geometry, or effort. **Fix:** Install capture. In descending order of durability: the changed default; the built object; the trained person; the altered floor plan; the decision log; the written record. Ostrom's design principles are, read as engineering, a list of eta-raising devices. Chapter 4 has the mechanisms. **Confusable with:** FM-12 (a Leaky Vessel that used to work), FM-19. **See:** §4.

FM-12 . The Ritual That Became a Meeting **Damages:** eta, by decay, under a preserved cadence. C is unchanged or rising. **Clinical:** A recurring event was designed with a capture mechanism. The cadence survived; the capture did not. What remains is a periodic forcing function driving a filter with a broken input stage — full cost, no conversion. This is the most common failure of *successful* practices, because the cadence is what everyone can see and the capture is what nobody was assigned. **Signature:** A meeting nobody can cancel and nobody can justify. The original artifact — log, board, checklist — exists, is out of date, and is not referenced. **Measurement:** Pull the artifact. Check its last-modified date against the event's cadence. If the gap exceeds about three cadence-periods, eta for that ritual is near zero. Confirm by asking three attendees to name the last decision the ritual produced. **Price:** The event's full C continues to be charged against h while its contribution to `eta S` approaches zero. In the ward diagnosis at the end of this chapter, a handover with intact cadence and degraded capture ran at **46.04%** of its own ceiling. **Fix:** Restore the capture step, or cancel the cadence. Do not keep the cadence for morale. A ritual whose artifact is dead is a C term with a good reputation. **Confusable with:** FM-11 (never had capture), FM-19. **See:** §2, §4.

FM-13 . The Import With No Socket **Damages:** eta for the imported category specifically; u is raised nominally and never realized. **Clinical:** The organization correctly diagnoses low u and correctly hires the outsider — and then provides no channel through which the outsider's

difference can reach a decision. The new category exists on the roster and not in the causal path. The u gain is real in the survey and zero in the output. **Signature:** A deliberate, well-justified hire from an unrelated lineage who has not changed any decision in two quarters. The outsider asked for input after the option set is closed. **Measurement:** Decision-provenance audit for the imported category: in the last ten consequential decisions, count the ones where the imported member's input entered before the option set was closed. Zero or one is diagnostic. **Price:** The ceiling gain from N = 4 to N = 5 is **0.2231** nats. If the fifth category has no socket, the realized gain is 0.0000 nats and the salary is pure C. **Fix:** Build the socket before the import: a named decision gate the new category owns, a veto scoped narrowly enough to be usable, a standing agenda slot that precedes option-closure. Chapter 26's sequencing rule applies — the socket is stabilization, and it precedes growth. **Confusable with:** FM-02 (which is an import with no *difference*; this is a real difference with no *channel*). **See:** §9, §26.

FM-14 . The Sprayed Seed **Damages:** eta for a diffusion intervention specifically. **Clinical:** A new practice is introduced by broad thin exposure — everyone hears once. Complex contagion (Centola & Macy 2007; Centola 2010) requires *multiple reinforcing contacts* before adoption, so thin broadcast reaches everyone and converts nobody. **Clustered topology beats random long ties**, the opposite of the reach intuition that produced the spraying. **Signature:** A launch with excellent awareness metrics and no adoption. Pilot sites chosen for geographic spread. "Everyone has been informed" as the implementation report. **Measurement:** Count *reinforcing* contacts, not contacts. For each non-adopter, how many of their frequent contacts have adopted? If the modal answer is one, the topology is wrong regardless of total exposure. **Price:** Adoption is threshold-driven and therefore discontinuous (Granovetter 1978); below threshold the return on exposure is near zero, which is why sprayed campaigns produce a flat curve rather than a slow one. **Fix:** Cluster first, do not spray. Saturate one dense neighbourhood to well above local threshold, then let the boundary do the work. Size the committed minority against Centola et al. 2018: roughly **25%** flipped an established convention in controlled online experiments — for a 22-household block that is 6 households, for a 12-person group it is 3 people. Badged: the 25% figure is from controlled online experiments and field transfer is a modelling assumption. **Confusable with:** FM-11. **See:** §19.

FM-15 . The Founder Battery **Damages:** Q is stored in a person rather than in the system; delta becomes a step function. **Clinical:** One individual is the conversion mechanism, the archive, and the maintenance regime. Measured eta is high and measured delta is low, and both readings are attributes of a person, not of an institution. The parameters are correct right up until the day they are catastrophically wrong. **Signature:** A named person who is in every consequential thread. Onboarding that consists of asking them. Nobody can state the escalation path without saying their name. **Measurement:** The absence test, run deliberately and at low stakes: the person takes two consecutive periods fully off-channel. Measure Q before and after. The decline is the person-attributable share of Q. Chapter 24 calls this the single most informative measurement a healthy organization can afford to run. **Price:** A commons block that lost its founding steward saw commons-units fall from 47 to 38 over nine months of near-zero input: `delta = 0.0233/month`, half-life **29.35 months**. The decay was gentle only because that block's stored order was physical. Where Q is procedural, the same test typically returns a far larger number. **Fix:** Externalize before

you replace. Stewards of record for every asset; the artifact carries the state, not the person (see FM-16's inverse and Chapter 21's stigmergy rule: **put the state in the medium, not in a coordinator**). **Confusable with:** FM-10, FM-19. **See:** §3, §21, §23.

Class III — Retention faults (delta and Q)

FM-16 . The Sealed Vault **Damages:** delta -> 0 with `S = 0`. Q is perfectly preserved and never updated. **Clinical:** The system has excellent retention and no input. From the exact solution (Theorem 3.1), with `S = 0` and `C = 0` the trajectory is `Q_t = (1 − delta)^t Q_0` — a pure decay so slow it reads as permanence. The vault is not failing; it is succeeding at the wrong thing. **A system with no forgetting cannot learn.** **Signature:** Procedures whose rationale nobody can reconstruct. Reverence for the archive combined with no additions to it. Change proposals answered with precedent, not argument. **Measurement:** Two readings. First, the input audit: over the last twelve periods, how many elements of Q were *added or amended*? Second, delta under zero input: if delta is low and additions are zero, the vault is sealed. **Price:** delta = 0.01 gives a half-life of **68.97** periods and a time constant of 100. delta = 0.005 gives a half-life of **138.28** periods. At those rates an obsolete rule outlives every person who might have questioned it. **Fix:** Raise delta deliberately in the compartments where you want learning: sunset clauses, mandatory re-justification cycles, term limits, planned demolition. Keep delta low where you want durability. **Delta is not a single number for a whole institution; it is a field over its compartments, and treating it as one number is itself the fault.** **Confusable with:** FM-18 (which is refusal to shed *at cost*; the vault costs nothing and simply ossifies). **See:** §1, §3, §23.

FM-17 . The Maintenance Trap **Damages:** h, via the `delta Q` term. No parameter is individually abnormal. **Clinical:** Q has grown to the point where `delta Q` alone consumes the ceiling. From Identity 3, `h = eta ln N − delta Q − C`; h is *decreasing in Q*. **Success consumes its own headroom.** This is the structural reason institutions ossify, and it is not a moral failure — it is an arithmetic one. **Signature:** An organization that was building five years ago and is now entirely maintaining, with no decision anyone can point to that caused the change. Every new initiative crowded out by obligations to existing ones. **Measurement:** Compute the stall point: `Q_stall = eta ln N / delta` with C = 0. Compare to current Q. If you are within about 30% of the stall point, the trap is already closing. **Price:** With eta = 0.40 units/nat and N = 12 (`eta ln N = 0.9940`): at delta = 0.02 the stall point is `Q = 49.70`; at delta = 0.05 it is **19.88**; at delta = 0.08 it is **12.43**; at delta = 0.12 it is **8.28**. A system with delta = 0.12 cannot hold more than about eight units of order no matter how well it interacts. **Fix:** Four levers and only four (Chapter 6): raise eta, raise N, cut C, or **accept lower Q**. The last is the one nobody considers, and shedding stored order to regain headroom is a legitimate and underused move — see FM-18 for the failure to make it. **Confusable with:** FM-18, FM-20. **See:** §6, §23.

FM-18 . The Unshed Q **Damages:** Q too high for the system's own delta; h negative as a consequence. **Clinical:** The organization retains stored order it no longer uses, cannot maintain, and will not retire. Every retained unit charges `delta` against headroom every period, forever. Unlike the Sealed Vault, the Unshed Q is *expensive*: the retained order actively consumes maintenance. **Signature:** An asset register nobody has pruned. Modules, documents, programs, or committees unused for a year that no one will delete. The phrase "someone might still need it." **Measurement:** Usage audit with a hard window. For every element of Q, when was it last used by someone other than its owner? Sum `delta * Q_unused`. That number is the annual rent on sentimentality, in the same units as your growth budget. **Price:** In Chapter 28's research group, retiring 12 of 46 assets moved `delta Q` from `0.1120 * 46 = 5.154` to `0.075 * 34 = 2.550` assets/quarter — a **50.5%** reduction in the maintenance charge, obtained by *removing* things. Shedding alone moved the one-step balance from **-3.387** to **-2.043** assets/quarter before any other lever was touched. **Fix:** Sunset by default. Every element of Q carries an expiry and must be re-justified to survive it. Archive tiers marked explicitly unmaintained are honest and cheap; pretending to maintain is neither. **Confusable with:** FM-17 (the trap is the arithmetic; the Unshed Q is the refusal to act on it). **See:** §6, §23, §26.

FM-19 . The Invisible Infrastructure **Damages:** delta, rising unobserved; then Q, discontinuously. **Clinical:** Star & Ruhleder (1996): infrastructure becomes visible only on breakdown. Weiser (1991): the most profound technologies are those that disappear. The consequence for this law is structural — **maintenance is chronically underfunded because working infrastructure produces no observable signal**, so delta rises for years with no reading on any instrument, and the first reading is a failure. **Signature:** No maintenance line in the budget. No named owner for a system everyone depends on. A backup nobody has restored from. The question "who maintains this?" producing a pause. **Measurement:** You cannot measure invisible infrastructure by watching it; you must *interrupt* it. Scheduled, announced, low-stakes interruption: take the dependency off-line for one period and measure the drop in Q. That number is the infrastructure's contribution, and it is the only honest basis for its budget. **Price:** Undetected delta drift is the most expensive error in the catalog because it compounds. A delta that drifts from 0.05 to 0.15 moves the half-life of stored order from **13.51** periods to **4.27** — a threefold acceleration that produces no symptom until the stock is already gone. **Fix:** Make the invisible visible on a schedule: interruption drills, maintenance dashboards that display *what did not break*, named stewards. Budget maintenance as a fixed fraction, not as a residual. **Confusable with:** FM-15, FM-12. **See:** §3, §23.

FM-20 . The Negative Fixed Point **Damages:** the model itself — `eta F(u*) < c*`, so `Q_inf = (eta F(u*) − c*) / delta < 0`. **Clinical:** Theorem 3.5 gives the steady state `Q_inf = (eta F(u*) − c*) / delta`. When conversion times realized surplus is below cost, the equilibrium is negative. **The theorem does not hide this failure mode and neither may you.** A negative fixed point means the regime is unsustainable or the linear model must be replaced — those are the only two readings, and the organization is usually reporting a third: "transitional." **Signature:** A multi-year "transition" with no defined end state. Reserves drawn down on a schedule described as temporary. **Measurement:** Compute `Q_inf` explicitly at every survey. It is one division. If it is negative or below current Q by a wide margin, the trajectory is heading there and the arrival date is computable from `tau = 1/delta`. **Price:** In Chapter 28's open-source case the measured pre-

intervention fixed point was `Q_inf = 20.39` modules against a current Q of 63 — the project was not stable, it was 43 modules above its own equilibrium and descending at **-2.47** modules/month, with a computed zero-crossing at **25.5 months**. **Fix:** Change a parameter or change the model. Both are legitimate; continuing to report the current regime as steady is not. If cost cannot fall below `eta F(u*)`, the system must shrink to a scale where it can, and shrinking on purpose is cheaper than shrinking by collapse. **Confusable with:** FM-17, FM-23. **See:** §5, §6, §26.

Class IV — Dissipation faults (C)

FM-21 . The Quadratic Committee **Damages:** C, quadratically, against a logarithmically growing ceiling. **Clinical:** Brooks (1975): communication paths grow as `n(n−1)/2`. The surplus ceiling grows as `ln N`. **Cost grows quadratically while surplus grows logarithmically**, so for any fixed per-channel cost there is a group size beyond which the group is strictly worse than a smaller one. Adding people to a late project makes it later. **Signature:** Meetings that require pre-meetings. Decisions requiring more sign-offs than arguments. A group that has grown and whose output has not. **Measurement:** Tally C directly in hours or in Q-units per period (Chapter 24's cost observable). Then compute the crossover: for per-channel cost `c0`, find the n where `c0 * n(n−1)/2` exceeds `eta ln n`. **Price:** Channels per nat of ceiling: n = 5 -> 10 channels, **6.21** per nat; n = 12 -> 66, **26.56**; n = 20 -> 190, **63.42**; n = 30 -> 435, **127.90**; n = 150 -> 11,175, **2,230.26**. With eta = 1.0 and per-channel cost 0.01, headroom peaks near n = 10 (`h = +1.853`), is still positive at n = 20 (`h = +1.096`), and **crosses zero at n = 27**. **Fix:** Nest, do not enlarge. Alexander's semilattice ("A City is Not a Tree," 1965) beats both the flat group and the strict tree: overlapping clusters preserve cross-contact while cutting channels. Chapter 28's commons case cuts 231 channels to **55.5** — **24.03%** — with four overlapping clusters of 5.5 households plus a delegate layer. **Confusable with:** FM-22 (which is about *load*, not *size*), FM-07. **See:** §5, §17, §21.

FM-22 . The Utilization Cliff **Damages:** C, non-linearly, as utilization approaches 1. **Clinical:** Standard queueing: for an M/M/1 queue, expected wait scales as `1/(1 − rho)`. Near full utilization the derivative explodes, so the last few percent of capacity are bought at enormous and accelerating cost in delay. **Slack is not waste; it is a cost-control device**, and removing it is the most common way an efficiency programme destroys an institution. **Signature:** Everything is scheduled and everything is late. Efficiency metrics at record highs alongside collapsing responsiveness. No one has unallocated time and nothing gets done. **Measurement:** Compute rho for the binding resource: arrival rate over service rate. For people, rho is the fraction of available hours already committed. Anything above about 0.85 sustained is in cliff territory. Little's Law (`L = lambda W`) closes the loop and lets you check your queue length reading against your wait reading. **Price:** Wait as a multiple of service time: rho = 0.50 -> **2.0x**; 0.80 -> **5.0x**; 0.90 -> **10.0x**; 0.95 -> **20.0x**; 0.98 -> **50.0x**. Moving a team from 80% to 95% booked multiplies waiting by **4.0x**; from 80% to 98% by **10.0x**. Queue length rises from 4.0 to 19.0 to 49.0 over the same range. **Fix:** Design to a target rho, typically 0.70-0.85 for work with variable service times. Protect slack explicitly and name it, so it cannot be reallocated by an efficiency review. Reduce variability

before raising utilization — variability, not utilization, is what makes the cliff steep. **Confusable with:** FM-21. **See:** §5.

FM-23 . The Unbooked Cost **Damages:** C, underestimated; h, overestimated; every downstream decision, poisoned. **Clinical:** The cost ledger records the costs that have invoices. Commute, emotional labour, context-switching, compliance, waiting, the repair of relationships damaged by the last decision — these dissipate stored order and appear on no line. A survey that omits them computes a headroom that does not exist and then spends it. **Signature:** Repeated overruns with no identifiable cause. People visibly exhausted by a workload the ledger describes as moderate. **Measurement:** Time-and-motion on the *whole* period, not on the tasks. Everything not converting S into Q is C. Chapter 24 defines the cost observable as a residual for exactly this reason: measure total capacity, measure productive capacity, take the difference — rather than enumerating cost items and hoping the list is complete. **Price:** In Chapter 28's research group, coordination ran at **17.4%** of 6,240 person-hours per quarter — 1,087 hours — which at the measured yield rate of **0.00147936** assets/person-hour priced C at **1.608** assets/quarter against a ceiling `eta ln N = 8.105`. Omitting it would have overstated h by **120%** of its true value of 1.343. **Fix:** Book C as a residual, not as a list. Then validate: the survey's predicted one-step change must reproduce the observed one. Chapter 28 shows all four cases reproducing observed decline to within 0.12 units; that agreement is what tells you the cost ledger is complete. **Confusable with:** FM-29 (both are survey faults; this one under-reports C, that one mis-specifies Q). **See:** §5, §24.

FM-24 . The Abrupt Threshold **Damages:** C, at every zone boundary; secondarily u, by suppressing crossing. **Clinical:** Two zones of different character meet with no transition. Every crossing costs a full reorientation: acoustic, social, visual, procedural. Alexander's pattern lineage treats the threshold as a device in its own right — impedance matching between zones — and its absence is a tax collected on every crossing, forever. **Signature:** A door people hesitate at. An open-plan floor where nobody uses the "collaboration zone" because entering it is a public act. **Measurement:** Count crossings per hour against the crossings the function requires, by direct observation (Whyte's method). Desire lines are the measured gradient; where the desire line does not match the designed line, the threshold is wrong. **Price:** Context-switching is a C term and it compounds with FM-22: at rho = 0.90 each additional interruption costs **10x** its own service time in downstream delay. **Fix:** Build the transition: porch, genkan, airlock, vestibule, buffer zone, staged announcement. In organizations the equivalent is the staged decision — a proposal that has a public-but-uncommitted phase. Depth, sightline into the next zone, and a place to pause are the three specifications. **Confusable with:** FM-26 (both make crossing costly; one is geometric, one is acoustic). **See:** §15.

Class V — Room and instrument faults

Separated because their fixes are capital purchases with known lead times, and because they are the cheapest interventions in this manual per unit of effect.

FM-25 . The Head of Table **Damages:** u, structurally and continuously. **Clinical:** A table with a head assigns a direction to the group. Everyone below the salt has a function, and that function is agreement. **Agreement between aligned instruments is echo, not evidence.** The geometry does this without anyone enforcing it, which is precisely why it is durable: furniture governs continuously and costs nothing to administer. **Signature:** Speaking order tracks seating order. Turn-taking Gini high. Contributions addressed to one seat rather than to the room. Sightlines that do not close. **Measurement:** Two instruments. (1) Turn-taking equality: record speaking time per participant for three meetings and compute the Gini coefficient — Woolley et al. (2010) found collective intelligence c correlates with *equality* of conversational turn-taking. (2) Sightline reciprocity: for all pairs, can each see the other's face without moving? Report the failing fraction. **Price:** The Parallel Wall prices the result. At N = 12, driving u from 0.30 to 0.05 by geometry alone moves S from **1.4586** to **0.4383** — a **70.0%** loss of realized surplus, purchased with a table shape. **Fix:** Round or polygonal geometry with full sightline reciprocity. Size the table by Brooks (n(n−1)/2 channels) and by acoustics jointly. Add protocol: silent-write-first, round-robin, last-speaks-first for rank. **Furniture is the cheapest governance instrument ever built. Confusable with:** FM-01, FM-26. The test: if the fault persists when the group meets elsewhere, it is not the table. **See:** §17, §18.

FM-26 . The Reverberant Room **Damages:** u, by suppressing the quiet voice; C, via the Lombard effect. **Clinical:** In a room with long reverberation, speech from previous syllables masks the current one. Signal-to-noise for the quiet speaker falls below intelligibility while the loud speaker remains intelligible. **The loudest speaker wins a reverberant room by physics, not by dominance** — and then the Lombard effect raises everyone's level, which raises the noise floor, which raises everyone's level again. **Signature:** People leaning in. Repetition requests. Quiet participants who are articulate elsewhere. Meetings that exhaust people disproportionately to their content. **Measurement:** Sabine: `RT60 = 0.161 V / A` (metric), where V is volume in m^3 and A is total absorption in m^2 sabins, `A = sum(S_i * alpha_i)` plus occupants at roughly 0.45 m^2 each at 500 Hz. Measure or compute both, occupied and empty. Add a background noise reading. **Price:** A 16.0 x 9.5 x 4.2 m reading room (`V = 638.40 m^3`) with hard finishes has `A = 40.750` and `RT60 = 2.5223 s` occupied, **3.1480 s** empty. Treated — 90 m^2 of ceiling absorber at alpha = 0.90, 40 m^2 of wall panel at 0.85, 100 m^2 of heavy carpet at 0.25 — A rises to **173.150** and RT60 falls to **0.5936 s**, a **4.25x** reduction. Ceiling treatment alone gives 0.8699 s; ceiling plus walls gives 0.6845 s. **Fix:** Absorption, in order of cost-effectiveness: ceiling first (largest uninterrupted surface, first reflection path), then the wall opposite the primary talker, then floor. **Acoustic design is c-factor engineering and it is the cheapest c-factor intervention available. Confusable with:** FM-25, FM-01. Acoustic faults are the ones that vanish when the same group meets in a different room — that is the discriminating test. **See:** §14.

FM-27 . The Single Setpoint **Damages:** u, by narrowing who can occupy the space; C, via avoidance. **Clinical:** The environment is optimized to one value — one temperature, one illuminance, one noise floor, one furniture height — chosen as the population mean. Individual tolerance varies widely across the six thermal variables (air temperature, radiant temperature, humidity, air

speed, clothing, metabolic rate) and across light sensitivity. An optimized setpoint is correct for the mean occupant and wrong for most actual occupants. **Adjustability outperforms optimization because it raises the range of people who can be present at all**, and presence is prior to every other parameter. **Signature:** Personal heaters and fans in a climate-controlled building. Blinds permanently down. Complaints in both directions about the same setting. **Measurement:** Occupancy mapping by seat over two weeks — Whyte's method applied indoors. Persistently unoccupied positions in an otherwise-used room are an environmental reading, not a furniture preference. Cross-check with a two-sided complaint tally. **Price:** Any occupant excluded is a category removed from N or a member removed from a category, priced by Identity 1. Excluding one of five staffed categories moves the ceiling from `ln 5 = 1.6094` to `ln 4 = 1.3863`, a **13.9%** loss — for a thermostat. **Fix:** Provide gradient and control rather than a single value: a range of thermal zones, task lighting alongside ambient, seating variety, sun and shade both available. Ulrich (*Science*, 1984) — post-surgical recovery faster with a window view of trees — is an argument for providing the good condition, not for averaging it away. **Confusable with:** FM-24. **See:** §16.

Class VI — Method faults (the engineer's own)

These damage nothing in the system and everything in your knowledge of it. Listed last; check them first, because a method fault invalidates every other reading.

FM-28 . Goodhart Capture **Damages:** the measurement, and then Q, once the measurement drives behaviour. **Clinical:** Goodhart's Law and Campbell's Law: a measure that becomes a target ceases to be a measure. Once the proxy Y for stored order is rewarded, effort flows to Y rather than to Q, and the correlation that justified the proxy is destroyed by the act of using it. **Signature:** A metric improving faster than any plausible underlying change. Improvement concentrated exactly at the measurement boundary. Practitioners who can explain how the number is produced but not what it means. **Measurement:** Hold out a second, unrewarded proxy for the same Q and track both. Divergence between rewarded and unrewarded proxies is the capture signal. Rotate the unrewarded proxy periodically so it does not become rewarded by attention. **Price:** Not directly computable in nats — the damage is that your h estimate is now fiction, so every subsequent number in this catalog is fiction. **Fix:** Never reward the proxy you use for measurement. Keep an audit proxy that is not published, not incentivized, and rotated. Prefer proxies that are expensive to fake and cheap to verify — physical state, third-party usability, independent re-execution. **Confusable with:** FM-29. **See:** §23, §25.

FM-29 . The Order Proxy That Was Really the Input **Damages:** the entire estimation. The fitted model is a tautology. **Clinical:** Paper E §6.3 names this explicitly and it is **the most common measurement error in the framework**. If the chosen order observable Y is a repackaging of the interaction observable U, then `Y_{t+1} = a Y_t + b F(U_t) − C + e` fits beautifully and says nothing. You have measured interaction twice and called the correlation a law. **Signature:** A proxy defined in terms of meetings, messages, commits, attendance, or engagement. Extremely

good fit on the first attempt. The proxy falls to zero the instant interaction pauses. **Measurement:** The lag test. Genuine stored order *persists* when input stops — that is what "stored" means, and it is why delta is measurable at all. Pause input for one to three periods and watch Y. If Y tracks the pause with no lag, Y is U. If Y decays with a measurable time constant, Y is Q, and you have just measured delta as a bonus. **Price:** Every parameter downstream is unidentified. There is no partial credit. **Fix:** Choose Y as a *state of the world outside the interaction*: modules that pass tests on the current toolchain; assets a non-author can independently re-run; patrons who return without being invited; shared capabilities usable without asking anyone. Chapter 28 applies this test explicitly to all four of its cases and rejects the naive proxy in each. **Confusable with:** FM-28. **See:** §24, §25.

FM-30 . The Unbadged Postulate **Damages:** the warrant. Nothing in the system; everything in the standing of your claims. **Clinical:** A modelling choice, an analogy, or a transfer from a laboratory result is presented with the same confidence as a theorem. The house rule from Book One: **the mysticism is permitted, but it is permitted in the open, with its badge on.** An unbadged postulate is not a lie; it is an unpriced liability, and it comes due when someone competent checks. **Signature:** No distinction in the prose between "Theorem 3.1 gives" and "we assume." A controlled online experiment cited as a field constant. The axiom S3 (affine effective count) invoked as though derived. **Measurement:** Read your own document with two highlighter colours: proven, and assumed. If the assumed colour does not appear, you have not found it; you have failed to look. **Price:** Nothing quantitative. A single discovered unbadged claim discredits the badged ones, which is a very poor trade for the confidence gained. **Fix:** Badge in the same sentence. "S3 is a design axiom — chosen, not derived." "The 25% figure is from controlled online experiments; field transfer is a modelling assumption." "The claim that this institution is described by a single scalar Q is a postulate." Appendix E collects every one of them for the whole volume. **Confusable with:** FM-31. **See:** CANON §2, §5; Appendix E.

FM-31 . The Borrowed Ceiling **Damages:** the plan. Requests `r_t > h_t`, which Corollary 4.2 forbids. **Clinical:** The design assumes synergy beyond additivity — that combining two sources of difference will yield more than their sum, or that a good enough team escapes its own headroom. Theorem 4.1 gives an upper bound of exactly the sum: `f(u) <= f(u_cross) + f(u_div)`. **No synergy beyond additivity is promised, and honest design does not assume one.** Separately, the triangle inequality fails for geometric triples (Paper A §5), so three-party reasoning cannot be chained metrically at all. **Signature:** A target that exceeds `r_max`. A plan justified by "these two things will multiply." Growth projections with no h calculation anywhere in the document. **Measurement:** Compute `h_t = eta ln N − delta Q_t − C_t` and compare to the requested `r_t`. Proposition 4.1: the increment is feasible only if `delta Q_t + C_t + r_t <= eta ln N`. This is one line of arithmetic and it is not optional. **Price:** At N = 8, `u_cross = 0.30` and `u_div = 0.25` give the honest bound `1.1314 <= 1.5790 <= 2.1430`. A plan assuming multiplication rather than the bound over-promises by up to **36%** relative to the true value even inside the permitted range. **Fix:** Plan to the bound. State `r_t <= h_t` as a hard constraint in the proposal document, and put the computed h next to the requested r on the same page. **Confusable with:** FM-07, FM-30. **See:** §6, §10, §26.

FM-32 . The Consent Bypass Damages: nothing in the law. It is a refusal, not a parameter. **Clinical:** An intervention raises a parameter by removing someone's ability to decline — mandatory disclosure of dissent, forced rotation, attendance that cannot be refused, "voluntary" participation with recorded non-participation. **CANON §5 places this outside the manual entirely**, and the reason is engineering as much as ethical: coerced u is not u. A person compelled to voice a difference voices the difference the compeller expects, which returns you to FM-01 by a longer road. **Signature:** A protocol with no opt-out. Non-participation that is visible and therefore costly. Consent obtained from a group rather than from each person. A benefit argued at the level of the collective whose cost falls on identifiable individuals. **Measurement:** For every protocol in the intervention design, answer one question in writing: *what happens to a person who declines?* If the answer includes any adverse consequence, the protocol is out of scope. **Price:** Not applicable. This is a boundary, not a trade-off, and it is not priced because it is not for sale. **Fix:** Redesign so the intervention works by removing resistance rather than by applying force. Chapter 20's method is the template — build the lower valley next door, and the worst case is that someone declines to move. Every intervention in Chapter 28 is constructed to satisfy this constraint, and where one could not be, it was cut. **Confusable with:** nothing. If you are unsure whether an intervention is a Consent Bypass, it is. **See:** §20, §22, §26; CANON §5.

The Specification — how to use the catalog

S27.1 — Enter by symptom, not by name. Do not scan the entry names for one that sounds like your situation; names are memorable and therefore misleading. Enter through Appendix D's symptom index, or through the parameter you can already measure.

S27.2 — Run the method faults first. Before any Class I-V diagnosis, clear FM-29 (order proxy), FM-23 (unbooked cost), FM-28 (Goodhart capture) and FM-30 (unbadged postulate). A method fault invalidates every other reading, and three of the four are cheap to check.

S27.3 — Measure the six parameters before naming anything. The survey (Chapter 24) is the entry condition for this chapter, not an alternative to it. Required: Q and its proxy Y; delta from decay under zero input; eta from attributable `delta_Q` per unit S; N by capability partition; u by two instruments (`u_cross` and `u_div`); C as a residual.

S27.4 — Compute both master numbers. Headroom `h = eta ln N − delta Q − C`, and realized fraction `S / ln N`. Record them together. §27.3 explains why either alone will mislead you.

S27.5 — Validate the survey against an observed change. Compute the predicted one-step change `eta S − delta Q − C` and compare it to the *measured* change in Q over a recent period. If they disagree by more than about 10% of the ceiling term, the survey is wrong and no diagnosis from it is admissible. This single check catches most instances of FM-23 and FM-29.

S27.6 — Expect three to five co-occurring faults, not one. Real systems present in clusters, and the clusters are structured: a Reverberant Room *causes* a Head of Table to bind harder; a Ceiling Chase *causes* a Quadratic Committee; a Founder Battery *hides* an Invisible Infrastructure. Name all of them, then rank.

S27.7 — Rank by lever cost, not by severity. Chapter 26's ranking, restated: eta is usually cheapest; C next; u next; N most expensive and least rewarding because it is logarithmic. Shedding Q is frequently the cheapest of all and is almost never considered.

S27.8 — If h < 0, you may not build. Corollary 3.4: if `C_t > eta ln N − delta Q_t`, no arrangement of people prevents order from decreasing. The permitted moves are: cut C, raise eta, shed Q. Growth is not on the list until h is positive.

S27.9 — Predict a number, and state what would falsify it. Every diagnosis terminates in a predicted trajectory with dated checkpoints and a written falsification condition. Chapter 25's four baselines — persistence, direct-input, matched autoregressive, and the shuffled-input control — are all four required. **A framework that cannot fail cannot be trusted.**

S27.10 — Apply the consent test to every proposed intervention. FM-32. In writing. Before costing.

The Failure Modes — of diagnosis itself

D-1 . Naming before measuring. The catalog's names are vivid, and vividness produces premature confidence. A named fault with no measurement behind it is a story. *Signature:* the diagnosis arrived within ten minutes of arriving on site. *Fix:* S27.3 is an entry condition.

D-2 . The single-fault reflex. Looking for *the* problem. Systems fail in clusters (S27.6), and fixing one member of a cluster while leaving the others returns the system to its previous state within a few time constants. *Fix:* require at least three named faults or an explicit written statement of why fewer suffice.

D-3 . Fixing the loudest fault. The most visible fault is rarely the binding one. In Chapter 28's library case the loud fault was conflict between patrons; the binding constraint was u, and the instrument was ceiling absorption. *Fix:* rank by effect on the one-step condition, computed, not by salience.

D-4 . The one-instance check. Verifying a fix at one site, one meeting, one week, and generalizing. A single successful instance is consistent with almost every hypothesis. *Fix:* pre-specify the number of instances and the checkpoints before running the pilot.

D-5 . Treating delta as a scalar for the whole system. Delta is a field over compartments. Documentation, relationships, physical plant, skills, and norms decay at wildly different rates, and an average delta is a number that describes none of them. *Fix:* estimate delta per compartment; report the vector.

D-6 . Diagnosing outside the model's range. The linear storage law has no saturation term. When predicted `Q_inf` exceeds a hard physical or organizational cap, the model is being used outside its range and must be replaced there. *Fix:* always state the cap alongside the fixed point, and say which one binds. The worked example below does exactly this.

D-7 . Mistaking the survey for the intervention. Measurement changes behaviour, usually temporarily and favourably. A survey period that shows improvement has told you about the survey. *Fix:* Chapter 25's shuffled-input control — the baseline most often omitted and most often decisive.

Worked Example — a hospital ward shift handover

A 34-bed inpatient ward. The complaint presented by the charge nurse is "handover is not working." Nothing else is offered. We take it through S27.1-S27.9.

The survey

Period: one week. **Order observable Y:** the number of currently-admitted patients for whom an incoming staff member can, unprompted, state the three current concerns and the escalation trigger. Measured by spot audit at three random times per week. `Q_0 = 21` of 34.

FM-29 check (S27.2). Is Y a repackaging of interaction? Run the lag test. Handovers were cancelled for three consecutive weeks during a staffing crisis. Y fell 27 -> 21 over those three weeks — a decay with a time constant, not a collapse to zero. Y is stored order, not input.

delta. From that same natural experiment, with `S ~ 0`:

```
delta = 1 - (21/27)^(1/3) = 0.080359 per week
half-life = ln 2 / (-ln(1 - delta)) = 8.274 weeks
```

N. Five categories are declared as handover participants: nursing, medicine, pharmacy, physiotherapy/OT, and discharge coordination. Two attend. **FM-06, the Frozen Denominator.** `N = 2`, `ln 2 = 0.693147` — not `ln 5 = 1.609438`. The ward has been planning against a ceiling **132% higher** than the one it owns.

u. Independent-concern-listing before discussion, 18 patients, two attending categories. Mean overlap of listed concerns, treated as the alignment cosine: `rho = 0.79`.

```
u = 1 - rho^2 = 1 - 0.6241 = 0.375900
N_eff = 1 + (2-1)(0.3759) = 1.375900
S = ln(1.375900) = 0.319108 nats     -> 46.04% of ln 2
```

eta. Over 12 weeks the ward logged 25 carried-patients attributable to specific handover exchanges, against cumulative `S = 12 x 0.319108 = 3.8293` nats.

```
eta = 25 / 3.8293 = 6.528614 carried-patients per nat
```

C. Cost observable as a residual (FM-23): productive handover time is 38 minutes/week; the audit finds 21 minutes/week of interruption, repetition, and downstream rework attributable to handover failure. Conversion rate `25 / (12 x 38) = 0.054825` carried-patients per productive minute.

```
C = 21 x 0.054825 = 1.151316 carried-patients per week
```

The two master numbers

```
h  = eta ln N - delta Q - C
   = 6.528614 x 0.693147 - 0.080359 x 21 - 1.151316
   = 4.525290 - 1.687531 - 1.151316
   = +1.686444 carried-patients per week          -> headroom is POSITIVE

S / ln N = 0.319108 / 0.693147 = 46.04%          -> realized fraction is LOW
```

Validation (S27.5). Predicted one-step change:

```
eta S − delta Q − C = 2.083333 − 1.687531 − 1.151316 = −0.755513 per week
```

Observed: Y fell 24 -> 21 over the preceding four weeks, or **-0.750** per week. Residual **-0.006**. The survey reproduces the observed decline and is admissible.

The diagnosis

Five co-occurring faults, ranked by effect on the one-step condition:

1. **FM-06, the Frozen Denominator.** Three of five declared categories do not attend. `N = 2`.
2. **FM-25, the Head of Table.** Handover is delivered by the outgoing senior nurse standing, to an audience seated. Speaking order tracks rank in all three observed sessions.
3. **FM-26, the Reverberant Room.** Handover happens in a corridor alcove, 4.2 x 2.6 x 3.0 m.

```
V = 32.760 m^3
A = 10.92 x 0.02 + 10.92 x 0.04 + 40.80 x 0.05 + 6 x 0.45
  = 0.2184 + 0.4368 + 2.040 + 2.70 = 5.395 m^2 sabins
RT60 = 0.161 x 32.760 / 5.395 = 0.9776 s
```

Nearly a full second in a room where six people must exchange precise clinical detail in under forty minutes, with corridor traffic passing through. 4. **FM-12, the Ritual That Became a Meeting.** The written handover sheet exists. Its template was last revised four years ago; the audit finds it filled in for 6 of 34 patients. 5. **FM-01, a partial Parallel Wall.** `u = 0.3759` between two categories whose training, supervision, and incentives substantially overlap.

The binding constraint is not h. Headroom is +1.69/week and the ward is losing 0.76/week. The ward is operating at 46% of a ceiling that is itself artificially low. This is precisely the case §27.3 warns about, and a diagnosis that stopped at “h is positive, we are fine” would have been exactly wrong.

The intervention

Ranked by lever cost per S27.7, and every element checked against FM-32 — none of the four requires anyone to participate who declines, and the pharmacy and discharge attendance is negotiated as a scheduling change rather than imposed.

#	Move	Lever	Target
1	Restore the handover sheet as a live board, filled during handover, not after	eta, delta	delta 0.0804 -> 0.065
2	Pharmacy and discharge coordination attend; physiotherapy does not (scheduling infeasible, stated honestly)	N	N 2 -> 4

#	Move	Lever	Target
3	Acoustic treatment of the alcove: 10.92 m^2 ceiling raft at alpha 0.90, 12 m^2 wall panel at 0.85	u, C	RT60 0.98 -> 0.22 s
4	Silent-write-first: each category lists concerns independently before the verbal round	u	rho 0.79 -> 0.68

Acoustics after treatment:

```
A = 0.2184 + 9.828 + (40.80 − 12) x 0.05 + 12 x 0.85 + 2.70 = 24.386 m^2 sabins
RT60 = 0.161 x 32.760 / 24.386 = 0.2163 s
```

Post-intervention parameters:

```
N = 4          ln N  = 1.386294
rho = 0.68 ->  u     = 0.537600
               N_eff = 2.612800
               S     = 0.960422 nats    -> 69.28% of ln 4
delta = 0.065  eta   = 6.528614 (held constant; conservative)
C = 26 x 0.054825 + 0.20 = 1.625450     (more attendees, plus board upkeep)
```

The prediction

```
h_new  = 6.528614 x 1.386294 − 0.065 x 21 − 1.625450
       = 9.050581 − 1.365000 − 1.625450 = +6.060131 per week

step   = eta S − delta Q − C
      = 6.270227 − 1.365000 − 1.625450 = +3.279777 carried-patients per week

Q_inf  = (6.270227 − 1.625450) / 0.065 = 71.46
tau    = 1/delta = 15.385 weeks
```

D-6 applies and must be stated. `Q_inf = 71.46` exceeds the ward census of 34. The linear law has no saturation term; beyond `Q = 34` the model is outside its range and must be replaced. **The binding limit is the census, not the fixed point.** The honest prediction is therefore a saturation date, not an equilibrium:

```
t = 1   Q = 24.280
t = 2   Q = 27.346
t = 4   Q = 32.895
t = 5   Q ~ 35.4  -> census cap reached; model invalid beyond this point
```

Predicted outcome: **the ward reaches full carried-patient coverage in four to five weeks and holds it.** Do-nothing comparison from the same start: Q = 18.32 at week 4, 16.41 at week 8,

14.76 at week 13.

What would falsify it

Preregistered before the first panel is hung, per S27.9 and Chapter 25:

1. **Trajectory.** Spot-audit Y at weeks 2, 4, and 8. At an eta 20% below the estimate, week 4 still predicts `Q = 28.35`. **If Y is below 28 at week 4, the eta estimate is refuted** and the conversion model requires re-specification.
2. **Acoustic mediation.** RT60 measured post-installation must fall below 0.35 s. If it does not, the absorption schedule was wrong and prediction 1 is void regardless of outcome.
3. **The u channel.** Independent-concern-listing repeated at week 6 must show `rho <= 0.72`. If u does not move while Y does, Y is rising for a reason outside the model — most likely the Hawthorne pattern named in D-7 — and the causal claim is refuted even though the outcome is favourable.
4. **Shuffled-input control.** Refit the minimal model with the interaction series randomly permuted across weeks. If the shuffled fit matches the ordered fit on held-out data, the apparent dependence on S is spurious. **This is the baseline most often omitted and most often decisive, and its omission would invalidate the whole diagnosis.**
5. **Persistence and matched-AR baselines.** The fitted model must beat both on held-out weeks. If a plain autoregression on Y alone predicts as well, the framework has added nothing here and should be reported as having added nothing.

Any one of 1, 3, 4, or 5 failing refutes the diagnosis. Failure of 2 refutes only the installation.

The catalog's job ends here. It named five faults, priced four of them, and produced a falsifiable number. Whether the number survives contact with the ward is the subject of Chapter 25 and the responsibility of whoever hangs the panels.

Chapter 28 . Worked Systems

Everything in this manual has been leading to four numbers computed four times.

The previous twenty-seven chapters established a law, six parameters, a gauge, a survey method, a falsification discipline, and a catalog of faults. None of that is worth anything if it cannot be carried through a real system from first observation to a dated prediction that could turn out to be wrong. This chapter does that four times, on four systems chosen because they have almost nothing in common: a twelve-person research group, a public library reading room, a distributed open-source project, and a residential commons of twenty-two households.

Every number below is computed, not asserted. Where a figure would require research that has not been done, the chapter says what would need measuring instead of supplying a plausible value. The four cases are constructed instruments — they are worked demonstrations of the method, not field reports — and that is stated here once so it never has to be qualified again. What is real in them is the arithmetic and the discipline: each survey is required to reproduce an independently observed change in Q before any diagnosis from it is admitted, and all four do, to within 0.12 units.

The four were chosen to span the diagnostic space. One has positive headroom and is shrinking anyway. One has enormous headroom and collects less than a third of it. Two have negative headroom and are forbidden to build. **No two of them have the same binding constraint, and in no case is the binding constraint the one the system's own members named.**

The Law

28.1 What a worked case must contain

Seven steps, in this order, no substitutions:

```
1. SURVEY          five operational choices + six parameters, measured
2. VALIDATE        predicted one-step change must reproduce the observed one
3. COMPUTE       h = eta ln N − delta Q − C, and the realized fraction S / ln N
4. BIND        identify which term is binding, by computed effect, not by salience
5. DESIGN       intervention, levers ranked by cost, consent test applied to each
6. PREDICT      a dated trajectory with a fixed point and a stated validity range
7. FALSIFY      what observation would refute it, written down before starting
```

Step 2 is the one most often skipped and it is the one that makes the rest admissible. Step 7 is the one that makes the exercise science rather than consulting.

28.2 The equations, restated for use

From Paper A (Theorem 2.1) and Paper E (§3, §6.4), quoted exactly:

```
S_t   = F(U_t) = ln(1 + (N−1)U_t)            surplus input, 0 <= S_t <= ln N
N_eff = 1 + (N−1)u                       effective count after redundancy audit
Q_{t+1} = (1 − delta)Q_t + eta S_t − C_t     the storage law
h_t   = eta ln N − delta Q_t − C_t           headroom (the master gauge)
Q_inf = (eta F(u*) − c*) / delta           steady state, delta > 0 (Thm 3.5)
```

```
tau   = 1 / delta                              time constant, in periods
growth  <=>  eta S_t > delta Q_t + C_t         one-step condition (Cor 3.2)
```

Badged, in the same sentence as it is used: axiom S3, the affine effective count, is a *design axiom* — chosen, not derived — and every `N_eff` below inherits that choice. Badged also: that a research group, a reading room, a code repository, and a housing block are each described by a single scalar Q is a modelling postulate, not a theorem. Papers A and E prove things about the law. They prove nothing about these four systems.

28.3 The estimator for u

All four cases estimate u the same way, so the method is stated once. Take an alignment measure rho between the members' or categories' contributions on a live question — correlation of signed error, overlap of independently listed concerns, rank correlation of independently stated priorities. Treat rho as the cosine between unit state vectors. Then, directly from `u = 1 - (x.y)^2`:

```
u = 1 - rho^2
```

Where the two-source decomposition (Theorem 4.1) is available, `u = u_cross + u_div` is estimated with two separate instruments and both numbers are reported, because a single-instrument diversity reading is Failure Mode FM-08.

28.4 The honesty clauses, live in every case below

- The surplus functional as deployed is a **lower bound** on the ensemble's information. Every predicted trajectory is a floor, not a forecast of the best possible outcome.
- The linear law has **no saturation term**. Where a computed `Q_inf` exceeds a hard physical or organizational cap, the model is outside its range and must be replaced there. This happens in two of the five cases in this chapter and it is flagged both times.
- **No thermodynamic negentropy law is proved.** Surplus is not identified with Shannon entropy.
- Every external result cited — Woolley, Centola, Ostrom, Sabine, Brooks, Ulrich — is real, and the *transfer* of that result to the system at hand is a modelling assumption, badged at the point of use.
- No intervention below removes anyone's ability to decline. Where a candidate intervention would have, it was cut, and the cut is recorded.

The Engineering — four systems

CASE A . A twelve-person research group

The system. A university-affiliated laboratory: one principal investigator, three postdoctoral researchers, six doctoral students, one research engineer, one coordinator. Funding stable. The complaint, as stated by the PI: *"We are busier than we have ever been and we are producing less."*

A.1 The survey The five operational choices required by Paper E §5.1:

Choice	Setting
State encoder	Each researcher's methodological position, from a 12-item forced-choice instrument
Interaction observable	Cross-person working sessions in which a live analytic decision was made
Order observable Y	Count of **independently re-runnable research assets**
Time scale	One quarter (13 weeks)
Cost observable	Person-hours not converting interaction into assets, taken as a residual

An asset counts toward Y only if someone **other than its author** has successfully re-run it in the trailing two quarters. `Q_0 = 46` assets.

FM-29 check — is Y a repackaging of U? No. Run the lag test: during a two-quarter funding hiatus with near-zero cross-person interaction, assets fell 52 -> 41. Y decayed with a time constant rather than tracking the pause instantaneously. Y is stored order.

delta. From that same hiatus, `S ~ 0`:

```
delta = 1 − (41/52)^(1/2) = 0.112046 per quarter
half−life = ln 2 / (−ln(1 − 0.112046)) = 5.833 quarters  (about 18 months)
```

N. Categories defined by capability, not by title — *what can this category independently detect, fix, or veto?* Four staffed: bench assay, imaging, statistics/ML, clinical translation. `ln 4 = 1.386294`.

u. Independent-estimate-before-discussion on 20 live analytic questions; mean pairwise correlation of signed error `rho = 0.86`. Seven of the twelve trained under the same advisor.

```
u     = 1 − 0.86^2 = 0.260400
N_eff = 1 + 3(0.2604) = 1.781200
S     = ln(1.781200) = 0.577287 nats     ->  41.64% of ln 4
```

Twelve people are purchasing the surplus of 1.78.

eta. Over eight quarters, 27 assets were attributable to specific logged working sessions against a cumulative `S = 8 x 0.577287 = 4.6183` nats.

```
eta = 27 / 4.6183 = 5.846309 assets per nat
```

C. Residual method (FM-23). Twelve people at 40 h/week x 13 weeks = 6,240 person-hours per quarter. Coordination — scheduling, reporting, compliance, internal review, meeting overhead — tallies at 17.4%, or 1,087 h/quarter. Over the eight estimation quarters the lab produced 61 durable assets from 41,234 productive person-hours:

```
yield = 61 / 41,234 = 0.00147936 assets per person−hour
C     = 1,087 x 0.00147936 = 1.608069 assets per quarter
```

A.2 The two master numbers, and validation

```
eta ln N  = 5.846309 x 1.386294 = 8.104705
delta Q   = 0.112046 x 46       = 5.154136
C         =                       1.608069

h = 8.104705 − 5.154136 − 1.608069 = +1.342500 assets/quarter
S / ln N  = 0.577287 / 1.386294   =  41.64%
```

Validation (step 2). Predicted one-step change:

```
eta S − delta Q − C = 3.375000 − 5.154136 − 1.608069 = −3.387205 assets/quarter
```

Observed: assets fell 60 -> 46 over the preceding four quarters, or **-3.500** per quarter. Residual **+0.113**, about 1.4% of the ceiling term. The survey is admissible.

A.3 The binding constraint Headroom is **positive** and the lab is losing three and a half assets a quarter. This is the case Chapter 27 §27.3 exists to prevent misreading: h is the ceiling, u decides how much of it you collect, and the lab collects 41.64%.

Break-even requires:

```
S* = (delta Q + C) / eta = 6.762205 / 5.846309 = 1.156662 nats
```

Against a ceiling of `ln 4 = 1.386294`, that is feasible — but it demands `N_eff = 3.179`, hence `u = 0.7264`, hence `rho = 0.523`. Moving twelve researchers from 86% error-correlation to 52% is a large ask and, alone, leaves no margin whatever. **Single-lever solutions all fail here**, and the arithmetic says so before anyone tries one:

Single lever	Resulting one-step change
Baseline	**-3.387**
u alone: 0.2604 -> 0.6156	-0.646
Shed Q alone: 46 -> 34	-2.043
delta alone: 0.1120 -> 0.075	-1.683
C alone: 1.608 -> 1.213	-2.992
N alone: 4 -> 5	-2.590

Every one is still negative. The binding constraint is **the joint burden of delta Q and a parallel roster** — FM-17 (Maintenance Trap) co-occurring with FM-01 (Parallel Wall), with FM-02 (Pseudo-Diversity) underneath: the lab is demographically varied and methodologically single.

Two further readings, both from the survey rather than from opinion.

The PI's own proposal was refused by arithmetic. The stated plan was to recruit three more doctoral students. Twelve people already carry `n(n−1)/2 = 66` communication channels; fifteen carry 105, a 59% increase in coordination surface. The ceiling gain, if all three landed in genuinely new categories — they would not have; all three were to be recruited from the dominant lineage — would be at most `eta(ln 7 − ln 4) = 5.846 x 0.5596 = 3.27` assets/quarter, against a coordination increment the lab could not price. Recruiting into `u = 0.26` adds redundancy at full salary. **This is FM-07, the Ceiling Chase, and it is the standard response of a laboratory that has correctly noticed it is falling behind.**

Where the surplus is actually being lost. Decompose u with two instruments (Theorem 4.1). The lab measures `u_cross = 0.2100` between its four method categories and `u_div = 0.0504` within them; they sum to the measured `u = 0.2604`. Bench assay and imaging disagree; two statisticians never do.

```
f(u_cross) = ln(1 + 3 x 0.2100) = 0.4886
f(u_div)   = ln(1 + 3 x 0.0504) = 0.1408
bound:  max = 0.4886  <=  f(u) = 0.5773  <=  sum = 0.6294        (Theorem 4.1)
```

Four fifths of the lab's realized surplus comes from crossing categories and one fifth from disagreement inside them. **The cross-disciplinary structure is real; the internal dissent is nearly absent.** Hiring more disciplines will not fix an organization in which each discipline has one permitted opinion, which is why A4.4 buys `u_div` — the independent-estimate protocol — before it buys the fifth category.

A.4 The intervention Four moves, ranked by lever cost per Chapter 26. Each was checked against FM-32 — what happens to a person who declines? — and one candidate was cut: mandatory rotation into a second methodological track, which had no opt-out and is therefore out of scope for this manual.

A4.1 — Shed Q (cheapest). Retire 12 assets that no non-author has re-run in four quarters to an explicitly unmaintained archive tier. `Q: 46 -> 34`.

A4.2 — Cut delta. Every surviving asset gets a named steward who is not its author, and a scheduled quarterly re-run. Target `delta: 0.1120 -> 0.075` (half-life 5.83 -> **8.891** quarters). The steward re-runs cost about 120 h/quarter and are booked as cost, not hidden.

A4.3 — Cut C. Coordination from 1,087 h/quarter to 700 h; adding the 120 h of steward work gives 820 h. `C = 820 x 0.00147936 = 1.213079`.

A4.4 — Raise u, and only then N. Two instruments: independent-estimate-first on every consequential analytic decision (this breaks cascades before they form, FM-03), and one recruitment into a fifth genuinely distinct category — a methodologist from an unrelated trade, hired **with a socket**: a named decision gate before option-closure, per FM-13. Target `rho: 0.86 -> 0.62`.

```
N = 5,  ln 5 = 1.609438
u     = 1 - 0.62^2 = 0.615600
N_eff = 1 + 4(0.6156) = 3.462400
S     = ln(3.462400) = 1.241962 nats   ->  77.17% of ln 5
```

eta is held at its measured 5.846309 — conservative, since the capture reforms should raise it.

A.5 The prediction

```
h_new = 5.846309 x 1.609438 - 0.075 x 34 - 1.213079
      = 9.409272 - 2.550000 - 1.213079 = +5.646193 assets/quarter

step  = 7.260894 - 2.550000 - 1.213079 = +3.497815 assets/quarter

Q_inf = (7.260894 - 1.213079) / 0.075 = 80.64 assets
tau   = 1/0.075 = 13.333 quarters;  half of the gap closes in 8.891 quarters
```

The lab starts from 34 (after shedding) and rebuilds:

Quarter	Q predicted	Do-nothing (from 46)
0	34.00	46.00
1	37.50	—
4	**46.49**	34.56
8	**55.64**	27.45
12	62.34	23.03
20	70.83	—
40	78.58	—

The lab returns to its starting stock of 46 assets in four quarters and passes it thereafter. The honest part of this prediction is its slowness: `tau` is more than three years, and the system is still 12 assets short of its fixed point after a decade. **A low delta is what makes stored order durable and it is exactly what makes rebuilding slow.**

A.6 What would falsify it Preregistered before the first asset is retired, per Chapter 25. All four baselines required.

1. **Trajectory.** At eta 20% below estimate, quarter 4 still predicts `Q = 41.31` and quarter 8 predicts `Q = 46.66`. **If Q < 41 at quarter 4, or < 46 at quarter 8, the eta estimate is refuted.** At eta 20% above, the same checkpoints are 51.68 and 64.63.
2. **The u channel.** Re-run the independent-estimate instrument at quarter 3. It must show `rho <= 0.70`. If Q rises while rho does not move, Q is rising for a reason outside the model and the causal claim fails even though the outcome is favourable.
3. **The delta channel.** Re-measure delta at quarter 6 by sampling assets with stewards absent for one quarter. It must be at or below 0.09. If not, A4.2 did not work and the fixed point is wrong.
4. **Shuffled-input control.** Refit `Y_{t+1} = a Y_t + b F(U_t) − C + e` with the interaction series permuted across quarters. **If the shuffled fit matches the ordered fit on held-out quarters, the apparent dependence on S is spurious and the entire diagnosis is void.**
5. **Persistence and matched-AR baselines.** The fitted model must beat both on held-out data, with `0 <= a <= 1` and `b > 0`. If a plain autoregression on Y predicts as well, report that the framework added nothing here.

Failure of 1, 2, 4, or 5 refutes the diagnosis.

CASE B . A public library reading room

The system. The main reading room of an urban branch library: 16.0 x 9.5 x 4.2 m, terrazzo floor, plaster ceiling, painted plaster walls, 38 m^2 of glazing, 44 m^2 of filled shelving, typical occupancy 18. A silence rule is posted and enforced. The complaint, as stated by the branch manager: *“We spend all our time policing noise, and our regulars are disappearing.”*

B.1 The survey

Choice	Setting
State encoder	Category of use, assigned by direct observation (Whyte's method)
Interaction observable	Interaction-minutes, classified within- or between-category
Order observable Y	**Regularity stock**: distinct patrons with >= 4 visits in the trailing 8 weeks
Time scale	One month (4 weeks)
Cost observable	Patron-facing staff hours diverted to mediation and complaint handling

FM-29 check. The obvious proxy — door count — is the *input*, not the order, and using it would be the single most common measurement error in the framework. Regularity stock is a habit that persists when the room is closed. Lag test below confirms it. `Q_0 = 84` regulars.

delta. An 11-week closure for asbestos remediation, `S = 0` throughout. Regularity stock fell 96 -> 61. Eleven weeks is 2.75 monthly periods:

```
delta = 1 - (61/96)^(1/2.75) = 0.152021 per month
half-life = 4.203 months
```

Y decayed with a measurable time constant across the closure rather than collapsing. Y is stored order, and the closure paid for the delta estimate.

N. Six categories of use, defined by capability and need: quiet solo reading; student group study; children and caregivers; job-seeking and computer use; daytime shelter and rest; community meeting. `ln 6 = 1.791759`. Three of the six are effectively excluded by the silence rule — noted as FM-06 in the making, and revisited in B.3.

u, both sources (FM-08 requires two instruments).

```
u_cross = 0.0375   from 9 of 240 observed interaction-minutes crossing category boundaries
u_div  = 0.0963   from a within-category purpose-dissimilarity survey of quiet readers
u       = 0.133800
N_eff   = 1 + 5(0.1338) = 1.669000
S       = ln(1.669000) = 0.512225 nats    ->  28.59% of ln 6
```

The room contains six kinds of person and generates the surplus of 1.67.

eta. Over a five-month bilingual story-hour pilot, 71 new regulars were attributable to programming in the room, against cumulative `S = 5 x 0.512225 = 2.5611` nats:

```
eta = 71 / 2.5611 = 27.722212 regulars per nat
```

C. The same pilot consumed 1,580 patron-facing staff hours, which gives the conversion rate:

```
conversion = 71 / 1,580 = 0.044937 regulars per patron-facing staff hour
C = 54 h/month (mediation + complaint admin) x 0.044937 = 2.426582 regulars/month
```

B.2 The two master numbers, and validation

```
eta ln N = 27.722212 x 1.791759 = 49.671535
delta Q  = 0.152021 x 84        = 12.769789
C        =                        2.426582

h = 49.671535 − 12.769789 − 2.426582 = +34.475164 regulars/month
S / ln N = 28.59%
```

Validation. Predicted step: `14.200000 − 12.769789 − 2.426582 = −0.996371` regulars/month. Observed: the register fell 96 -> 84 over the preceding twelve months, **-1.000** per month. Residual **+0.004**. Admissible.

B.3 The binding constraint Headroom is **+34.5 regulars per month** — enormous — and the room is losing one regular a month. The room's own fixed point at current parameters is `Q_inf = 77.45`, so it is not merely declining but descending toward an equilibrium below where it stands.

The binding constraint is **u**, and it is binding at 28.59% of ceiling. The reason is physical. The silence rule exists because the room cannot acoustically hold two activities at once, and the silence rule is what drives `u_cross` to 0.0375. **The manager named noise as the problem and noise is the symptom; the room's reverberation time is the cause, and u is the parameter it damages.** This is FM-26 (Reverberant Room) generating FM-01 (Parallel Wall) through a posted rule.

Sabine, computed. `RT60 = 0.161 V / A` (metric), `V = 16.0 x 9.5 x 4.2 = 638.40 m^3`. Absorption coefficients at 500 Hz; occupants at 0.45 m^2 each.

```
A = 152.0 x 0.02  (terrazzo floor)      =   3.04
  + 152.0 x 0.04  (plaster ceiling)     =   6.08
  + 176.2 x 0.05  (painted plaster)     =   8.81
  +  38.0 x 0.04  (glazing)             =   1.52
  +  44.0 x 0.30  (filled shelving)     =  13.20
  +  18   x 0.45  (occupants)           =   8.10
                                        --------
                                   A    =  40.750 m^2 sabins

RT60 = 0.161 x 638.40 / 40.750 = 2.5223 s occupied
                                 3.1480 s empty
```

Two and a half seconds. In a room of this volume, speech from one syllable masks the next; the quiet voice falls below intelligibility while the loud voice does not; the Lombard effect then raises the floor for everyone. **The loudest patron wins this room by physics.** The staff are not policing behaviour; they are absorbing, in labour, an acoustic deficiency in the building.

B.4 The intervention **B4.1 — Absorption.** 90 m^2 of Class A ceiling absorber (alpha = 0.90) over the central span; 40 m^2 of wall panel (alpha = 0.85) on the wall opposite the primary reading bay; 100 m^2 of heavy carpet with underlay (alpha = 0.25) over the terrazzo.

```
A = 52.0 x 0.02 + 100.0 x 0.25 + 62.0 x 0.04 + 90.0 x 0.90
  + 136.2 x 0.05 + 40.0 x 0.85 + 38.0 x 0.04 + 44.0 x 0.30 + 18 x 0.45
```

```
  = 1.04 + 25.00 + 2.48 + 81.00 + 6.81 + 34.00 + 1.52 + 13.20 + 8.10
  = 173.150 m^2 sabins

RT60 = 0.161 x 638.40 / 173.150 = 0.5936 s          (4.25x reduction)
```

Partial packages, for budgeting: ceiling only gives **0.8699 s**; ceiling plus walls gives **0.6845 s**. Ceiling first is correct — largest uninterrupted surface, first reflection path, and it alone brings the room inside the usable band. The absolute cost is not stated here because it is local and this manual does not invent figures; what is specified is the absorption schedule, which is the deliverable a contractor prices.

B4.2 — Zoning, not silence. Replace the blanket rule with three acoustically distinct bays — a treated quiet bay, a group bay, a story alcove — arranged so that categories are in *sightline* of each other but not in each other's acoustic field. This is Alexander's nested-scales pattern and Gehl's edge effect used deliberately: people occupy perimeters and thresholds first, and a threshold that lets you see the next zone before you enter it is the device (FM-24).

B4.3 — A visible medium. A board on which patrons post and claim recurring uses of the room. Stigmergy, per Grassé 1959 and Chapter 21's rule: **put the state in the medium, not in a coordinator.**

B4.4 — The observable that tells you it is working before Q moves. Gehl's distinction between necessary, optional, and social activity supplies a leading indicator. Necessary activity — returning a book, using a terminal because there is nowhere else — happens regardless of quality. Optional and social activity happen only when conditions permit them. **Optional activity is the observable proxy for eta**, and it responds within weeks, where regularity stock takes months. Instrument it Whyte's way, by direct observation on a fixed schedule: count people who stay without a task, count conversations lasting over two minutes, and map seat occupancy. Whyte's findings about small urban spaces transfer directly to the interior — sittable space, sun, and triangulation are the variables — and the transfer is badged as a modelling assumption, since his observations were of plazas, not reading rooms.

Three specific readings to take before and after, none of which requires a budget:

Instrument	Pre-treatment reading	Post-treatment target
Occupancy by seat, 12 sessions	Persistently empty seats mapped	Empty-seat count halved
Conversations > 2 minutes / hour	Baseline count	Any increase; the silence rule made this near zero
Stay-without-task headcount	Baseline count	Increase, as optional activity returns

Two conditions are also worth stating because they are cheap and the literature on them is real: daylight and view. Ulrich (*Science*, 1984) found post-surgical recovery faster with a window view of trees. That is a hospital finding and its transfer to a reading room is a modelling assumption — but the 38 m^2 of existing glazing is an asset already paid for, and the seating plan should put the quiet bay against it rather than against the interior wall, at no cost.

Post-intervention parameters, all conservative:

```
u_cross: 0.0375 -> 0.19    (zoning admits cross-category contact at the edges)
u_div:   0.0963 -> 0.13      (more purposes admitted within each category)
u = 0.320000;  N_eff = 2.600000;  S = 0.955511 nats  ->  53.33% of ln 6
eta held at 27.722212        (no credit taken for the board)
delta held at 0.152021       (no credit taken for improved habit formation)
C = 22 h x 0.044937 + 1.20 (panel upkeep + a named room steward) = 2.188608
```

B.5 The prediction

```
h_new = 49.671535 - 12.769789 - 2.188608 = +34.713139 regulars/month
step  = 26.488890 - 12.769789 - 2.188608 = +11.530494 regulars/month
Q_inf = (26.488890 - 2.188608) / 0.152021 = 159.85 regulars
tau   = 6.578 months;  95% of the gap closes in about 19.7 months
```

Note carefully: **h barely moved — from +34.48 to +34.71 — while the one-step change moved from -1.00 to +11.53.** Headroom is almost untouched because it does not contain u at all. The entire intervention is a change in how much of an already-available ceiling the room collects. This is the clearest demonstration in the book of §27.3.

Month	Q predicted	Do-nothing
0	84.0	84.0
3	113.6	81.6
6	**131.6**	79.9
12	**149.4**	78.4
24	158.4	77.6
36	159.6	77.5

B.6 What would falsify it

1. **RT60.** Measured occupied reverberation time after installation must be **below 0.80 s**. If it is not, the absorption schedule was wrong and every downstream prediction is void independently of what happens to attendance.
2. **The u channel, both sources.** Re-run the interaction-minute classification at month 3. `u_cross` must exceed 0.12. If Q rises while `u_cross` does not, the mechanism claimed here is not the mechanism operating.
3. **Trajectory, banded by realized u.** The prediction is explicitly conditional:

Realized u	S	Q at month 6	Q at month 12	Q_inf
0.1338 (no change)	0.5122	80.9	79.7	79.0
0.20	0.6931	101.6	108.1	112.0
0.26	0.8329	117.6	130.1	137.5
0.32 (target)	**0.9555**	**131.6**	**149.4**	**159.8**
0.40	1.0986	148.0	171.9	185.9

If Q is below 101 at month 6, u did not move at all and the zoning failed.

4. **Mediation hours.** Must fall below 30 h/month by month 3. If reverberation falls and mediation does not, the conflict was not acoustic and C was mis-attributed.
5. **Shuffled-input control**, and the persistence and matched-AR baselines. As in Case A.

CASE C . A distributed open-source project

The system. A library of 141 modules, three core maintainers, 37 contributors active in the trailing year, no employer, no roadmap authority. The complaint, from the maintainers: *"We are drowning and the docs are a lie."*

C.1 The survey

Choice	Setting
State encoder	Contributor capability category — what they can independently fix
Interaction observable	Independent reviews and design discussions on merged changes
Order observable Y	Modules that are **CI-green on the current toolchain and whose documented examples execute**
Time scale	One month
Cost observable	Green modules destroyed by in-period breaking changes

FM-29 check, and a rejected first attempt. The first proposed proxy was "modules with a maintainer who responded within 14 days." That is maintainer *interaction* wearing a state costume, and it fails the test. The adopted proxy is a pure state of the artifact: it can be evaluated by a stranger with a clean machine and no access to the project's social graph. `Q_0 = 63` of 141 modules.

delta. An eight-week holiday freeze with near-zero merges. Green modules fell 71 -> 63, purely from toolchain drift, dependency deprecation, and link rot:

```
delta = 1 - (63/71)^(1/2) = 0.058021 per month
half-life = 11.596 months
```

N — and FM-06 in the open. Seven categories are declared in the contributor guide: core runtime, packaging, docs, platform ports, downstream integrators, security, accessibility/i18n. Three of the seven — ports, security, accessibility — have no active member.

```
declared: ln 7 = 1.945910
staffed:  ln 4 = 1.386294
```

The project has been planning against a ceiling 40.4% higher than the one it owns.

u. Forty pull requests were each reviewed independently by two reviewers who did not see each other's comments. Mean Jaccard overlap of flagged defects `rho = 0.68`:

```
u     = 1 - 0.68^2 = 0.537600
N_eff = 1 + 3(0.5376) = 2.612800
S     = ln(2.612800) = 0.960422 nats    ->  69.28% of ln 4
```

eta. Over ten months, 41 modules were newly greened or restored to green attributably to logged review and design threads, against `S = 10 x 0.960422 = 9.6042` nats:

```
eta = 41 / 9.6042 = 4.268955 modules per nat
```

C. Distinguished from delta by cause: delta is decay *without* input (measured in the freeze); C is damage caused *by the activity itself*. Over twelve months, 35 green modules were destroyed by in-period breaking changes and dependency bumps that were not repaired within the period:

```
C = 35 / 12 = 2.916667 modules per month
```

C.2 The two master numbers, and validation

```
eta ln N = 4.268955 x 1.386294 = 5.918028
delta Q  = 0.058021 x 63       = 3.655339
C        =                      2.916667

h = 5.918028 - 3.655339 - 2.916667 = -0.653978 modules/month     <-- NEGATIVE
S / ln N = 69.28%
```

Validation. Predicted step: `4.100000 - 3.655339 - 2.916667 = -2.472006` modules/month. Observed: green modules fell 92 -> 63 over the preceding twelve months, **-2.417** per month. Residual **-0.055**, under 1% of the ceiling term. Admissible.

C.3 The binding constraint `h < 0`. Corollary 3.4 states it without qualification: if `C_t > eta ln N - delta Q_t`, **no choice of interaction geometry prevents order from decreasing.** More reviewers will not save this project. Better reviewers will not save it. A larger community will make it worse, because adding contributors adds channels quadratically (Brooks) against a ceiling that grows logarithmically — FM-07, the Ceiling Chase, which is what most projects in this state attempt.

The measured fixed point makes the position explicit (FM-20):

```
Q_inf = (4.100000 - 2.916667) / 0.058021 = 20.39 modules
```

The project stands at 63 and its own equilibrium is 20. At the current rate it reaches zero green modules in **25.5 months**. Note the realized fraction: **69.28%** of ceiling — the *geometry is comparatively healthy*. The reviewers are genuinely decorrelated. That is precisely why the geometry is not the fix.

Binding: **C**, with a Frozen Denominator (FM-06) suppressing the ceiling and an Unshed Q (FM-18) inflating the maintenance charge. Also present and worth naming: three maintainers is a hub, and scale-free networks are robust to random failure and **fragile to targeted hub attack** (Albert, Jeong & Barabási 2000). The project's bus factor is its single largest unpriced risk, and it is FM-15, the Founder Battery, distributed across three people.

C.4 The intervention Per Chapter 26: with `h < 0` you may not build. The permitted moves are cut C, raise eta, shed Q. N is added here only after the first three, and only because two of the empty categories can be staffed from existing downstream users at near-zero coordination cost.

C4.1 — Cut C (the binding term). A compatibility CI matrix pinned to the current and two prior toolchains; a mandatory deprecation window; no merge without a downstream smoke test. This is stigmergy applied to code: the constraint lives in the medium, enforced by the machine, not in a maintainer's attention. Target `C: 2.9167 -> 1.10 modules/month`.

C4.2 — Shed Q. Move 22 modules that no downstream consumer has imported in twelve months to an explicitly unmaintained archive tier. Not deleted; **honestly labelled**, which is the entire point. `Q: 63 -> 41`, and `delta Q: 3.655 -> 2.379`.

C4.3 — Raise eta. A merge checklist that will not close without the durable artifact: an executable example, a docs entry, and a named module owner. Measured eta is 4.269; target **5.0** — deliberately below the 5.6 the pilot suggested, because a prediction should be made at the conservative end of its own evidence.

C4.4 — Raise N, with sockets. Staff two of the three empty categories — platform ports and security — from existing downstream users who already do this work privately. `N: 4 -> 6`, `ln 6 = 1.791759`. Each new category gets a named review gate it owns (FM-13). Accessibility is left empty and **recorded as empty**, because nobody has volunteered and inventing a category would be FM-06 all over again.

Adding a security reviewer and a platform porter to the review pool should also decorrelate review. Target `rho: 0.68 -> 0.52`:

```
u     = 1 - 0.52^2 = 0.729600
N_eff = 1 + 5(0.7296) = 4.648000
S     = ln(4.648000) = 1.536437 nats    ->  85.75% of ln 6
```

Ostrom's design principles, read as engineering, supply the governance layer: clear boundaries (the archive tier), congruence with local conditions (the toolchain matrix), collective-choice arrangements (contributors set the deprecation window), monitoring (CI), graduated sanctions (revert, then require a smoke test, then require co-review), conflict-resolution mechanisms, and recognized rights to organize. Badged: Ostrom's principles are an empirical generalization from common-pool resource institutions; their transfer to a code repository is a modelling assumption.

C4.5 — How to land the discipline, which is a contagion problem, not an announcement. A merge checklist and a deprecation window are behaviours, and behaviours spread by complex contagion: adoption needs multiple reinforcing contacts, so **clustered topology beats broad thin exposure** (Centola & Macy 2007; Centola 2010). Announcing the policy on the mailing list is FM-14, the Sprayed Seed, and it is what most projects do. Instead, saturate one dense neighbourhood of the contributor graph — the six people who review each other's work weekly — until the practice is unremarkable there, then let the boundary carry it. Sizing: Centola et al. (*Science*, 2018) found a committed minority near **25%** flipped an established convention; at 37 active contributors that is **10 people**. Badged: the 25% figure is from controlled online experiments, and field transfer to a volunteer code project is a modelling assumption, not a measurement.

Conway's Law is also load-bearing here and it cuts both ways. Systems mirror the communication structure of the organizations that build them (Conway 1967). The three empty categories are not an accident of recruitment: the library has no port layer and no security surface *because*

nobody in the communication structure owns those concerns. Staffing the categories will, over time, change the artifact's seams. That is the mechanism by which C4.4 eventually reduces C as well as raising N — and it is slow, so no credit for it is taken in the prediction below.

C.5 The prediction

```
h_new = 5.0 x 1.791759 − 0.058021 x 41 − 1.10
      = 8.958797 − 2.378872 − 1.100000 = +5.479926 modules/month

step  = 7.682185 − 2.378872 − 1.100000 = +4.203313 modules/month

Q_inf = (7.682185 − 1.100000) / 0.058021 = 113.44 modules
tau   = 17.235 months
```

Validity range, stated (FM-D6). After archiving, the project has 119 maintainable modules. `Q_inf = 113.44` sits just under that cap, and the linear law has no saturation term, so **the model is trustworthy only while Q is well below 119.** It crosses Q = 100 at month 29; beyond roughly Q = 95 the prediction should be treated as directional, not numerical.

Month	Q predicted	Do-nothing (from 63)
0	41.0	63.0
6	**62.8**	50.2
12	**78.1**	41.2
24	96.2	30.5
36	105.0	—

The project passes its pre-intervention stock of 63 green modules in month 6 — from a *smaller* base, because shedding was one of the levers.

C.6 What would falsify it

1. **The C channel first.** Count green modules destroyed by in-period breaking changes over months 1-6. **If the rate is above 2.0/month, C4.1 failed** and the following bands apply instead:

Scenario	h	step	Q at 12	Q_inf
Full package (eta 5.0, C 1.10)	+5.480	+4.203	78.1	113.4
eta unchanged (4.269), C 1.10	+4.170	+3.080	68.2	94.1
eta 5.0, C slips to 2.00	+4.580	+3.303	70.1	97.9
eta 5.0, C discipline fails entirely	+3.663	+2.387	62.1	82.1

   Every row is positive: **the intervention's sign is robust even to complete failure of its largest single component**, because shedding Q and staffing N together moved h from -0.65 to +3.66 before eta or C were touched. That robustness is itself a falsifiable claim — if Q is still falling at month 6 under any of these rows, the model is wrong, not the execution.

2. **The N channel.** Both new categories must land at least three independent reviews in the first quarter. A category with a name and no output is FM-06 restored, and the ceiling reverts to `ln 4`.
3. **The u channel.** Re-run the paired-review overlap measurement at month 6. `rho <= 0.60` required.
4. **Shuffled-input control** and the persistence and matched-AR baselines, as before. This case is the most vulnerable of the four to a spurious fit, because CI-green counts move for toolchain reasons entirely unrelated to human interaction. **If the shuffled control fits as well, the whole causal story is an artifact of the release calendar.**

CASE D . A residential commons / co-housing block

The system. Twenty-two households, 51 residents, purpose-built co-housing: common house with kitchen and laundry, guest room, workshop, tool library, shared garden beds, bike shed. Governed by whole-group consensus. The complaint, from a resident: *"Everything takes a year and nothing stays fixed."*

D.1 The survey

Choice	Setting
State encoder	Household activity category
Interaction observable	Work parties and commons decisions with cross-household participation
Order observable Y	**Shared capabilities usable without negotiation**, in commons-units (cu)
Time scale	One month
Cost observable	Household-hours consumed by governance, mediation, and free-rider repair

A capability counts toward Y if a resident can use it *without asking anyone*: it works, it is findable, and its rules are known. Walk-through inspection: 14 tool-library items, 7 common-house systems, 12 written norms with observed compliance at or above 80%, 5 residents able to run a work party unaided. `Q_0 = 38 cu`.

FM-29 check. Meetings held would be the input. Capabilities that work when nobody is meeting is the order. The lag test is available and decisive, below.

delta — and FM-15 in the field. After the founding steward moved out, work parties lapsed for nine months (`S ~ 0`). Commons stock fell 47 -> 38:

```
delta = 1 - (38/47)^(1/9) = 0.023341 per month
half-life = 29.348 months  (2.45 years)
```

This is the lowest delta in the chapter, for the obvious reason that most of this commons is **physical**. Bricks forget slowly.

N. Five household activity categories by need and contribution: families with young children; retired residents; full-time-employed households; remote workers; households with mobility limits. ln 5 = 1.609438.

u. Each household independently ranked twelve candidate commons projects before any discussion. Mean pairwise Spearman rank correlation rho = 0.41:

```
u     = 1 − 0.41^2 = 0.831900
N_eff = 1 + 4(0.8319) = 4.327600
S     = ln(4.327600) = 1.465013 nats    ->  91.03% of ln 5
```

This is the highest realized fraction in the chapter. The block is genuinely, measurably heterogeneous, and it converts almost none of that into anything.

eta. Over the fourteen months when monthly work parties ran, 19 cu of new or restored capability were attributable, against S = 14 x 1.465013 = 20.5102 nats:

```
eta = 19 / 20.5102 = 0.926369 cu per nat
```

C. Conversion rate from the same window: 19 cu from 1,344 productive household-hours of work party = 0.01413690 cu/h. Governance tally: one whole-group consensus meeting per month at 22 households x 2.0 h = 44 h; mediation 11 h; free-rider repair 7 h. Total 62 household-hours/month.

```
C = 62 x 0.01413690 = 0.876488 cu per month
```

D.2 The two master numbers, and validation

```
eta ln N = 0.926369 x 1.609438 = 1.490934
delta Q  = 0.023341 x 38       = 0.886966
C        =                       0.876488

h = 1.490934 − 0.886966 − 0.876488 = −0.272521 cu/month        <-- NEGATIVE
S / ln N = 91.03%
```

Validation. Predicted step: 1.357143 − 0.886966 − 0.876488 = −0.406311 cu/month. Observed: the walk-through five months ago read 40 cu; today it reads 38, or **-0.400** per month. Residual **-0.006**. Admissible.

D.3 The binding constraint h < 0 again — and for the exact opposite reason to Case C. Here the geometry is superb: 91% of ceiling, an honest 0.83 non-parallelism, five real categories. What the block has is **eta = 0.926 cu per nat** — the lowest conversion efficiency in the chapter by a factor of four.

The residents are having genuinely different conversations and almost nothing durable comes out of them. That is FM-11, the Leaky Vessel, in its pure form: full surplus, no capture. The work parties happen, the ideas are real, and the tool comes back unlabelled to a shelf nobody owns.

Second binding term: C. Twenty-two households in whole-group consensus is n(n−1)/2 = 231 communication channels — FM-21, the Quadratic Committee. At n = 22 the ratio is **74.73 chan-**

nels per nat of ceiling; at n = 5 it is 6.21. The block pays twelve times the coordination cost per unit of ceiling that a five-household cluster would.

Third: FM-15, and it is already historical — the block's delta was measured *because* one person left and everything stopped.

D.4 The intervention **D4.1 — Raise eta (the binding lever, and the cheapest).** Three capture devices, in ascending durability:

- **A steward of record for every capability.** Named, on the object, rotating annually, with the right to decline the rotation.
- **A physical commons board** in the common house: what exists, where it is, who stewards it, what is broken. Stigmergy — **the state lives in the medium, not in a coordinator** — which is also the structural correction to the Founder Battery.
- **A work-party close-out rule:** no party ends without one written durable artifact. A label, a repaired item logged, a norm written down. Chapter 4's finding restated: *"we had a great work party"* is an eta = 0 report.

Pilot-measured eta reached 1.55; the prediction is made at **1.40**, conservatively.

D4.2 — Cut C by nesting, not by shrinking. Four clusters of about 5.5 households, with **overlapping membership** — each household sits in two clusters. This is Alexander's semilattice ("A City is Not a Tree," 1965) rather than a tree, and the overlap is what preserves cross-contact.

```
whole group:   22 x 21 / 2                        = 231 channels
nested:         4 clusters x (5.5 x 4.5 / 2) = 49.5
             + delegate council 4 x 3 / 2   =  6.0
                                             -------
                                           55.5 channels     = 24.03% of 231
```

Whole-group consensus moves to quarterly and covers constitutional matters only; clusters decide monthly within scope. **Badged as a modelling assumption:** the projection below scales meeting hours by the channel ratio, on the reasoning that the binding cost is pairwise reconciliation rather than seat time. That assumption is itself in the falsification list.

```
meeting term:   44 h x 0.2403                    = 10.571 h
whole-group quarterly, amortized: 22 x 2.5 / 3   = 18.333 h
mediation (down from 11)                         =  8 h
free-rider repair (graduated sanctions, Ostrom 5) =  4 h
                                                   -------
                                                    40.905 h/month

C = 40.905 x 0.01413690 = 0.578267 cu/month
```

D4.3 — Ostrom's eight principles as the eta and C specification. This is the case the principles were derived for — a common-pool resource governed by its users — and each principle maps to a specific parameter in the law. Badged: Ostrom's principles are an empirical generalization from field studies of common-pool resource institutions (*Governing the Commons*, 1990); their statement is hers, the parameter mapping is this manual's.

Ostrom principle	Instrument in this block	Parameter
1. Clearly defined boundaries	Which capabilities are commons, which are household	eta, C
2. Congruence of rules with local conditions	Cluster-level rules for cluster-level assets	C
3. Collective-choice arrangements	Those affected set the rule; opt-out preserved	u, and CANON §5
4. Monitoring	The commons board shows broken items, by whom, since when	eta, delta
5. Graduated sanctions	Reminder -> repair request -> stewardship reallocated	C
6. Conflict-resolution mechanisms	A standing, low-cost, non-escalating process	C
7. Recognition of rights to organize	Clusters may form and dissolve without approval	u, N
8. Nested enterprises	The four-cluster semilattice itself	C

Principle 8 is the one this case turns on, and it is the same structural claim as Alexander's semilattice from a different literature: **nested, overlapping enterprise beats both the flat assembly and the strict hierarchy**, and the arithmetic in D4.2 prices exactly how much. Principle 4 is the one that does the eta work — a commons board is a monitoring device, and monitoring is what converts a conversation into a state of record.

Consent check (FM-32), and one cut. Graduated sanctions must preserve the right to decline. The version adopted escalates from a reminder, to a request for repair, to reallocation of the capability's stewardship — none of which removes anyone's ability to withdraw from the commons activity entirely. A proposed version that made cluster attendance a condition of tool-library access **was cut**, because it made non-participation costly and is therefore outside this manual's scope.

Honest cost of the nesting: u falls. Clustering, even with overlap, reduces cross-category contact. Predict `u: 0.8319 -> 0.79`:

```
N_eff = 1 + 4(0.79) = 4.160000
S    = ln(4.160000) = 1.425515 nats    ->  88.57% of ln 5   (down from 91.03%)
```

delta is held at 0.023341 — no credit taken for the board, though it should help.

D.5 The prediction

```
h_new = 1.40 x 1.609438 - 0.023341 x 38 - 0.578267
      = 2.253213 - 0.886966 - 0.578267 = +0.787980 cu/month

step  = 1.995721 - 0.886966 - 0.578267 = +0.530488 cu/month
```

```
Q_inf = (1.995721 - 0.578267) / 0.023341 = 60.73 cu
tau   = 42.843 months  (3.57 years);  half the gap closes in 29.3 months
```

Month	Q predicted	Do-nothing
0	38.0	38.0
12	**43.6**	33.7
24	**47.8**	30.5
36	51.0	28.0
60	55.2	—
120	59.4	—

This is the slowest system in the book and the arithmetic says so plainly. `3 tau` is 128.5 months — **ten years and eight months** to settle within 5% of the fixed point. The same low delta that let the commons survive nine months of neglect is what makes it take a decade to build. A prospectus promising visible transformation in a year would be lying, and the lie would be detectable from `tau = 1/delta` alone.

Single-lever check, to show the bundle is necessary:

Lever	step	h
C cut alone (eta 0.926)	-0.145	+0.026
eta alone (1.40, C unchanged)	+0.288	+0.490
Both	**+0.530**	**+0.788**

D.6 What would falsify it

1. **The eta channel, which is the whole claim.** Count attributable cu per nat over months 1-12. Bands:

eta realized	step	Q at 12	Q at 24	Q_inf
0.926 (no change)	-0.145	36.5	35.3	31.8
1.10	+0.103	39.1	39.9	42.4
1.40 (target)	**+0.530**	**43.6**	**47.8**	**60.7**
1.55	+0.744	45.9	51.8	69.9

   If Q is at or below 39 cu at month 12, eta did not move and the capture devices failed.

2. **The channel-scaling assumption.** Measure actual governance household-hours at month 6. The prediction requires **at or below 45 h/month**. If nesting cuts channels but not hours, the scaling assumption is refuted and C must be re-estimated from the measurement.

3. **The u cost.** Re-run the independent-ranking instrument at month 9. If `rho` rises above 0.55 — that is, u falls below 0.70 — the clusters have become a tree rather than a semilattice, and the overlap must be increased. This is the intervention's own most likely failure and it is predicted in advance rather than explained afterward.

4. **Persistence, direct-input, matched-AR, and shuffled-input controls.** All four. With a `tau` of 43 months, a persistence baseline is a genuinely strong competitor over any horizon shorter than two years, and **if the model cannot beat persistence at month 24 it has added nothing.**

The Specification

The general procedure, abstracted from the four cases. Numbered for checking.

S28.1 — Choose the five operational observables before measuring anything: state encoder, interaction observable, **order observable Y**, time scale, cost observable. Write them down. Changing one later invalidates every earlier reading.

S28.2 — Apply the lag test to Y before proceeding. Find or create a period of near-zero input and watch Y. If Y tracks the pause with no lag, Y is the input and must be replaced. **All four cases above rejected an obvious proxy at this step, and Case C rejected its first choice explicitly.**

S28.3 — Estimate delta from that same zero-input period: `delta = 1 - (Q_1/Q_0)^(1/t)`. The interruption that annoys you is the measurement you cannot otherwise afford.

S28.4 — Partition N by capability, not by title. Count only staffed categories. Record empty categories as empty.

S28.5 — Estimate u with two instruments, `u_cross` and `u_div`, and report both. `u = 1 - rho^2`. Report the realized fraction `S / ln N` in the same breath.

S28.6 — Estimate eta as `attributable delta_Q / cumulative S` over at least five periods.

S28.7 — Take C as a residual, not as a list: total capacity minus productive capacity, priced at the same yield rate used for eta.

S28.8 — **Validate.** Predicted `eta S - delta Q - C` against observed change in Q. Agreement within about 10% of the ceiling term. **If the survey does not reproduce the past, it may not be used to predict the future.**

S28.9 — Compute h and `S / ln N`. Together. Then:

```
h < 0   -> you may not build. Cut C, raise eta, or shed Q. In that order of cheapness.
h = 0    -> maintenance only. Hold.
h > 0    -> growth available. Any requested r_t must satisfy r_t <= h_t.
             If S / ln N is also low, the binding constraint is u, not h.
```

S28.10 — Rank levers by cost: eta cheapest, then C, then shedding Q, then u, then N. **N is last because it is logarithmic**: going 4 -> 5 buys 0.2231 nats of ceiling for a 25% increase in categories, and every added category adds channels quadratically.

S28.11 — Apply the consent test to every element: *what happens to a person who declines?* If the answer contains an adverse consequence, cut the element and record the cut.

S28.12 — Predict a dated trajectory, a fixed point, a time constant, and **a validity range**. State the cap. Where `Q_inf` exceeds a physical or organizational cap, say which one binds.

S28.13 — Write the falsification conditions before starting. Four baselines: persistence, direct-input, matched autoregressive, **shuffled-input**. Band the prediction by the uncertainty in your least reliable parameter — usually eta.

The Failure Modes

What went wrong across these four surveys, and what to expect in yours.

The named complaint was wrong every time. The PI said “busy and unproductive” and the fault was a maintenance charge plus a parallel roster. The branch manager said “noise” and the fault was 2.52 seconds of reverberation collapsing u. The maintainers said “drowning” and the fault was churn cost against a frozen denominator. The resident said “slow” and the fault was a conversion efficiency of 0.93. **The complaint reliably names the symptom with the highest salience, and salience is uncorrelated with bindingness.**

Positive headroom read as health (Cases A and B). Both had `h > 0` and both were shrinking. Any diagnosis stopping at h would have been exactly wrong in half the chapter. `h` is the ceiling; `S / ln N` is the collection rate; you need both numbers or neither.

Single-lever thinking (all four). Case A’s table shows every single lever leaving the step negative. Case D’s shows the same. Case C is the exception that proves the rule — its bundle is robust to failure of its largest component, and that robustness was *computed*, not hoped for.

Taking credit for unmeasured improvement. Every case above held at least one parameter fixed that the intervention should plausibly improve — Case B held both eta and delta, Case A held eta, Case D held delta. **A prediction made at the optimistic end of its own evidence is not a prediction.**

Model range violations (Cases C and E). `Q_inf = 113.44` against 119 maintainable modules; `Q_inf = 16.40` against a desk that holds perhaps a dozen threads. The linear law has no saturation term and will happily predict past the cap. Say where it stops being true.

The shuffled-input control is the one that would kill these. Case C is the clearest: green module counts move on the release calendar for reasons entirely unrelated to human interaction. Every other baseline can be passed by a model that is merely tracking a trend. **A framework that cannot fail cannot be trusted, and this is precisely the test these four cases would fail if the framework were empty.**

The catalog entries encountered, by case

Cross-referenced to Chapter 27, so the four surveys can be read as a worked index. Note that no case presented with fewer than three co-occurring faults, which is Chapter 27’s specification S27.6 confirmed rather than assumed.

Case	Faults found	Binding	Lever pulled first
A . Research group	FM-17 Maintenance Trap, FM-01 Parallel Wall, FM-02 Pseudo-Diversity, FM-18 Unshed Q, FM-07 Ceiling Chase (proposed, refused)	delta-Q with u	Shed Q
B . Library room	FM-26 Reverberant Room, FM-01 Parallel Wall, FM-06 Frozen Denominator, FM-24 Abrupt Threshold	u	Absorption
C . Open source	FM-06 Frozen Denominator, FM-20 Negative Fixed Point, FM-18 Unshed Q, FM-15 Founder Battery (x3), FM-14 Sprayed Seed (proposed, refused)	C	Cut C
D . Co-housing	FM-11 Leaky Vessel, FM-21 Quadratic Committee, FM-15 Founder Battery (historical)	eta	Capture devices
E . One desk	FM-11 Leaky Vessel, FM-19 Invisible Infrastructure	eta with delta	The index card

Two patterns are worth naming. **FM-11, the Leaky Vessel, and FM-15, the Founder Battery, appear in three cases each** — conversion failure and person-held storage are the two most common faults in this sample, and both are cheap to fix relative to everything else. And **the fault the system proposed to fix was, in two of four cases, a fault the arithmetic refused**: the research group's recruitment plan and the open-source project's mailing-list announcement were both stopped by a computation, not by an argument.

Worked Example — one desk, and the scale-free claim

The four cases span three orders of magnitude in headcount. This last one has a headcount of one, and it is here to test CANON §1.4 directly: the claim that the storage law is **scale-free in form**, so that only `(delta, eta, N, C)` change between a desk, a room, a team, a building, an institution, and a city. That is a claim about **model form**, not about physics, and it is badged as such wherever it appears.

The system. One person, one desk, three concurrent workstreams. Period: one week. Order observable: **resumable threads** — work whose state can be recovered from the desk alone, without reconstruction from memory. `Q_0 = 6`.

```
delta:  after a two-week absence, resumable threads fell 7 -> 5
        delta = 1 - (5/7)^(1/2) = 0.154846 per week;  half-life 4.120 weeks

N     = 3 workstreams;   ln 3 = 1.098612
u:      independent recall test: how often does resuming one thread require
        reconstructing another?  coupling rho = 0.74
        u = 1 - 0.74^2 = 0.452400
        N_eff = 1 + 2(0.4524) = 1.904800
        S = ln(1.904800) = 0.644377 nats   ->  58.65% of ln 3

eta:     13 threads made resumable over 10 weeks, attributable to end-of-
session capture
        eta = 13 / (10 x 0.644377) = 2.017452 threads per nat
yield:  13 / (10 weeks x 5.0 productive hours) = 0.260000 threads per hour
C:      2.39 h/week of re-orientation and searching
        C = 2.39 x 0.26 = 0.621400 threads per week

h = 2.017452 x 1.098612 - 0.154846 x 6 - 0.621400
  = 2.216398 - 0.929074 - 0.621400 = +0.665924 threads/week

step = 1.300000 - 0.929074 - 0.621400 = -0.250474 threads/week
```

Observed: threads fell 7 -> 6 over four weeks, **-0.250** per week. Residual **-0.0005**. The survey validates at a headcount of one, with the same equation, the same seven steps, and the same validation rule as a twenty-two household commons.

The intervention is one index card per thread, written at the end of each session and parked in view: what was decided, what is next, what would unblock it. It is the smallest possible instance of the same three devices used in Case D — a state of record, a visible medium, and a close-out rule.

```
delta: 0.154846 -> 0.09      (the card survives the week; memory does not)
rho:   0.74 -> 0.55  ->  u = 0.697500,  N_eff = 2.395000
S      = ln(2.395000) = 0.873383 nats   ->  79.50% of ln 3
C:     1.1 h x 0.26 = 0.286000

h_new = 2.216398 - 0.540000 - 0.286000 = +1.390398 threads/week
step  = 1.762009 - 0.540000 - 0.286000 = +0.936009 threads/week
Q_inf = (1.762009 - 0.286000) / 0.09 = 16.40 threads
tau   = 11.111 weeks
```

Week	Q predicted	Do-nothing
0	6.00	6.00
2	7.79	—
4	**9.27**	5.21
8	11.51	4.80

Week	Q predicted	Do-nothing
12	13.05	4.60

Validity range, again. `Q_inf = 16.40` exceeds what one person can hold; the practical cap is nearer ten to twelve. The model is trustworthy to about week 6 and directional after that. Naming the cap is not a weakness of the method; **failing to name it is the failure**, and it is Failure Mode D-6.

The cross-case comparison

Five systems, one equation, parameters spanning three orders of magnitude:

System	Period	Units	delta	eta	N	u	S/lnN	C	Q	h	step
Desk	week	threads	0.1548	2.017	3	0.452	58.7%	0.621	6	+0.666	-0.250
Research group	quarter	assets	0.1120	5.846	4	0.260	41.6%	1.608	46	+1.343	-3.387
Library room	month	regulars	0.1520	27.722	6	0.134	28.6%	2.427	84	+34.475	-0.996
Open source	month	modules	0.0580	4.269	4	0.538	69.3%	2.917	63	-0.654	-2.472
Co-housing	month	cu	0.0233	0.926	5	0.832	91.0%	0.876	38	-0.273	-0.406

After intervention:

System	delta	eta	N	u	S/lnN	C	h	step	Q_inf	tau
Desk	0.0900	2.017	3	0.698	79.5%	0.286	+1.390	+0.936	16.4	11.1 wk
Research group	0.0750	5.846	5	0.616	77.2%	1.213	+5.646	+3.498	80.6	13.3 q
Library room	0.1520	27.722	6	0.320	53.3%	2.189	+34.713	+11.530	159.8	6.6 mo
Open source	0.0580	5.000	6	0.730	85.8%	1.100	+5.480	+4.203	113.4	17.2 mo
Co-housing	0.0233	1.400	5	0.790	88.6%	0.578	+0.788	+0.530	60.7	42.8 mo

Five readings from the table, each of which is a finding rather than a restatement:

1. Every one of the five was losing order at survey, and no two for the same reason. Binding constraints were, respectively: delta and u jointly; delta-Q with a parallel roster; u alone; C with a frozen denominator; eta with a quadratic committee. A method that always finds the same fault is not diagnosing.

2. `h` and `S / ln N` are independent readings and both are required. The library has 53x the headroom of the co-housing block and collects less than a third of its ceiling. The block collects 91% of a ceiling it cannot afford. Either number alone points at the wrong lever.

3. eta spans a factor of thirty across the five systems and is the cheapest thing to change in four of them. This is Chapter 4's central claim, arrived at from the arithmetic rather than asserted.

4. tau = 1/delta sets the honest timescale, and it is not negotiable by effort. The library settles in about twenty months; the co-housing block needs a decade. **The same low delta that makes a commons resilient is what makes it slow, and no intervention in this manual changes that trade.**

5. The form held at every scale. One desk and twenty-two households used the same seven steps, the same validation rule, and the same falsification discipline, and both surveys reproduced their observed change to within 0.006 units. That is what the hyperrecursive symmetry claim in CANON §1.4 actually asserts and all it asserts: **same equation, different parameters, and the parameters must be measured every time.**

The measurements above are the point. Not the model — the model is four lines. **What this manual supplies is the discipline of putting a number where an opinion was, validating it against something that already happened, and writing down in advance what would prove it wrong.**